
Knowledge Integration to Support Decision Making

Alyson Wilson, Ph. D. (agw@lanl.gov)

Deborah Leishman, Ph.D. (leishman@lanl.gov)

C. Shane Reese, Ph.D. (reese@statmail.byu.edu)

April 16, 2002



Los Alamos 1945



Los Alamos 2002



LANL Statistical Sciences Group

Mission: Bring statistical reasoning and rigor to multi-disciplinary scientific investigations through development, application, and communication of cutting-edge statistical sciences research.

Statistical Sciences Focus Areas

- Reliability
- Information Integration Technology (IIT)
- Computer Model Evaluation
- Statistical Population Bounding
- Monte Carlo Methods
- Computational Statistics
- Biological Sciences Applications



What is IIT?

IIT is a combination of processes, methods, and tools for collecting, organizing, and analyzing diverse information from dynamic environments to support decision making under uncertainty.

- IIT brings together the data, information, and **distributed** knowledge of different scientific disciplines, organizational levels, and geographically separate teams.
- IIT makes advanced problem-solving **capability** and **defensibility** available to decision makers.



Goals of IIT

GOAL: Develop a “**standard**” framework of processes, methods, and tools useful for evolving R&D to support of decision making under uncertainty.

CURRENT PRACTICE: Data, modeling, and analysis has evolved in a stovepipe manner within disciplines. Integration of the science either occurs through some “test” event or in the mind of the decision maker.



IIT Approach

Create a framework for **integrating scientific knowledge**, to accelerate R&D, that is:

- **flexible** allowing all diverse and heterogeneous sources of information to be included
- mathematically **rigorous** and **traceable** to ensure confidence in the predictions
- **complete** and able to support dependent objectives
- **builds on the best** of what is already being done



Lesson Learned

The Problem is not Modeling, it is Decision Making

Optimal decision-making requires diversity of information:

- **Sources of information** - theoretical models, test data, computer simulations, expertise and expert judgment (from scientists, field personnel, decision-makers...)
- **Content of the information** - information about system structure and behavior, decision-maker constraints, options, and preferences...)
- **Multiple communities** that are stakeholders in the decision process



Partners in IIT Development



F-22 SPO/Seek Eagle



Up Front
GONE FISSION

DR. SPOCK MEETS
DR. STRANGELOVE

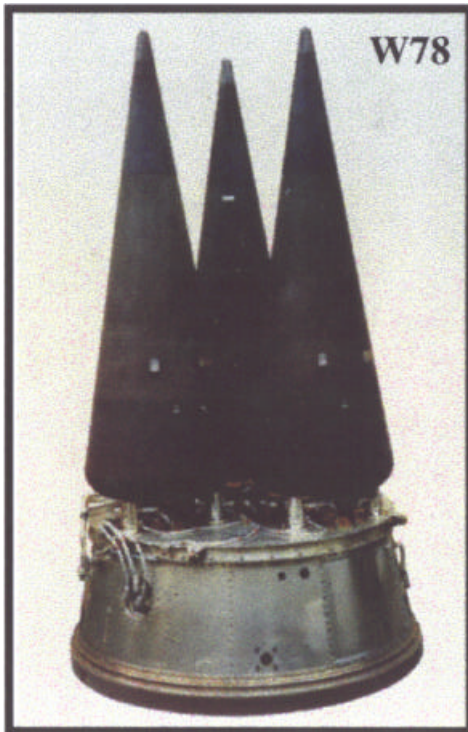


AMCOM/RDEC
Los Alamos



BUSINESS WEEK / July 10, 2000

Los Alamos Nuclear Weapons

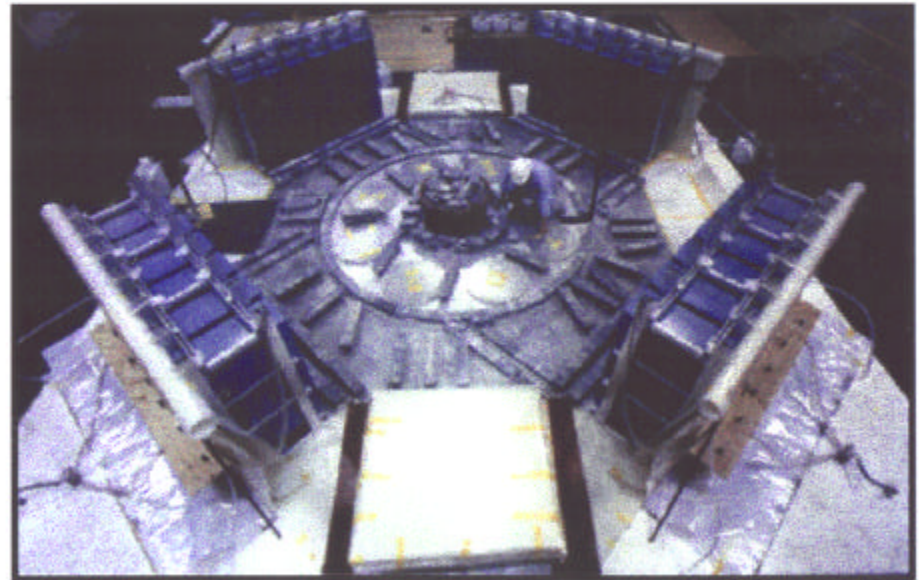


How We Carry Out Our Mission

Before: Design-Test-Produce



Now: Surveil-Assess-Respond



IIT Experience

Continuous and Comprehensive Evaluation of the “System”

- **Building confidence** in “system” performance, reliability, sustainability, dependability, etc
- **Resource allocation** (experimental design) and analysis for sub- and full “system” tests
- **Data/information requirements** for “system” assessment
- **Value of all information sources** including
 - Data on similar systems
 - Computer/simulation models
 - Experience/expertise, i.e., human judgment
 - Test data

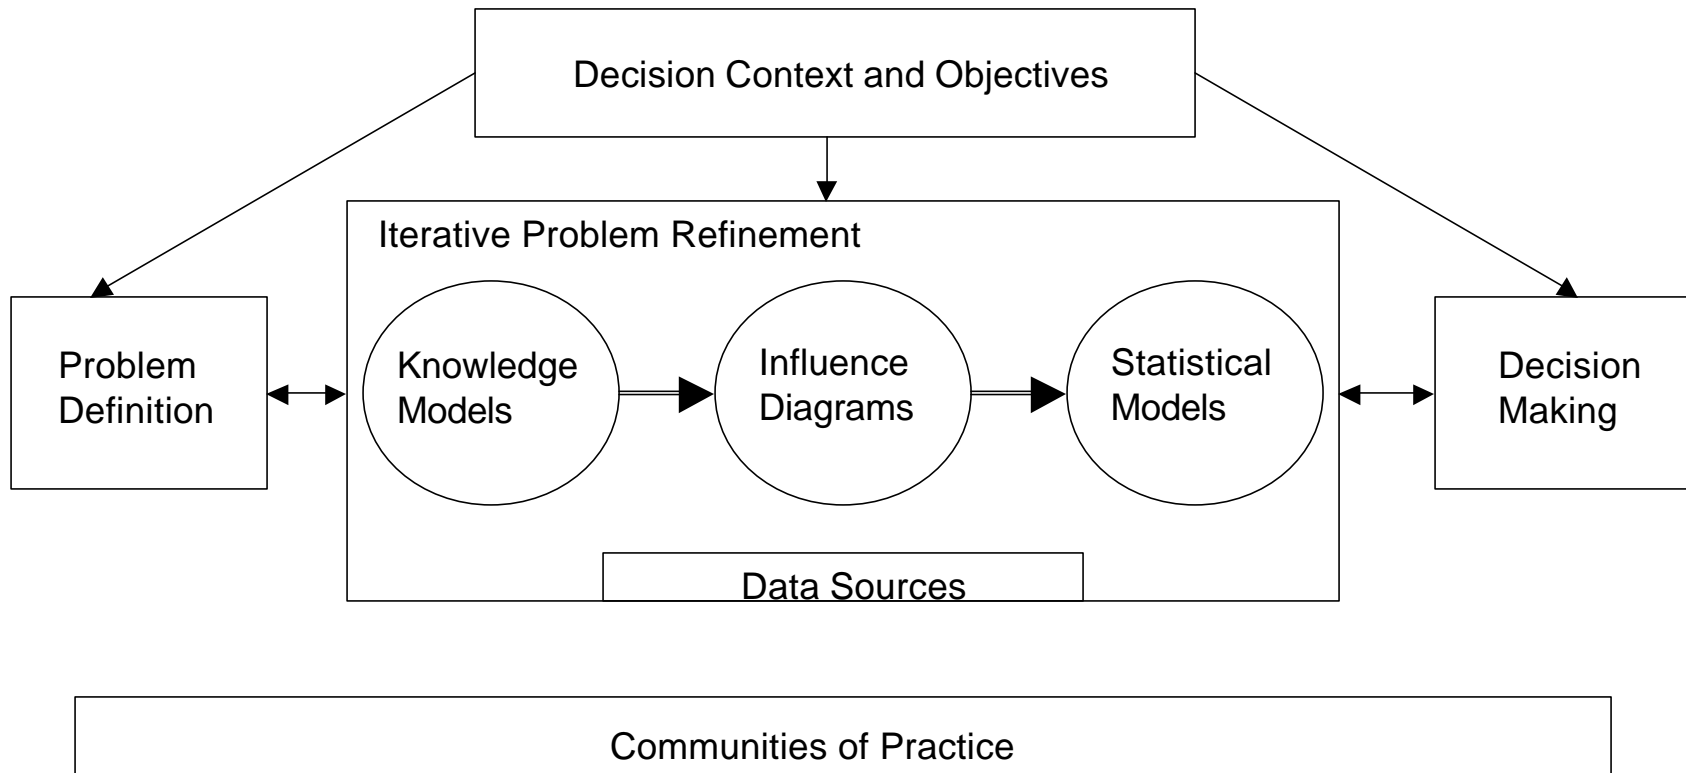


IIT Components

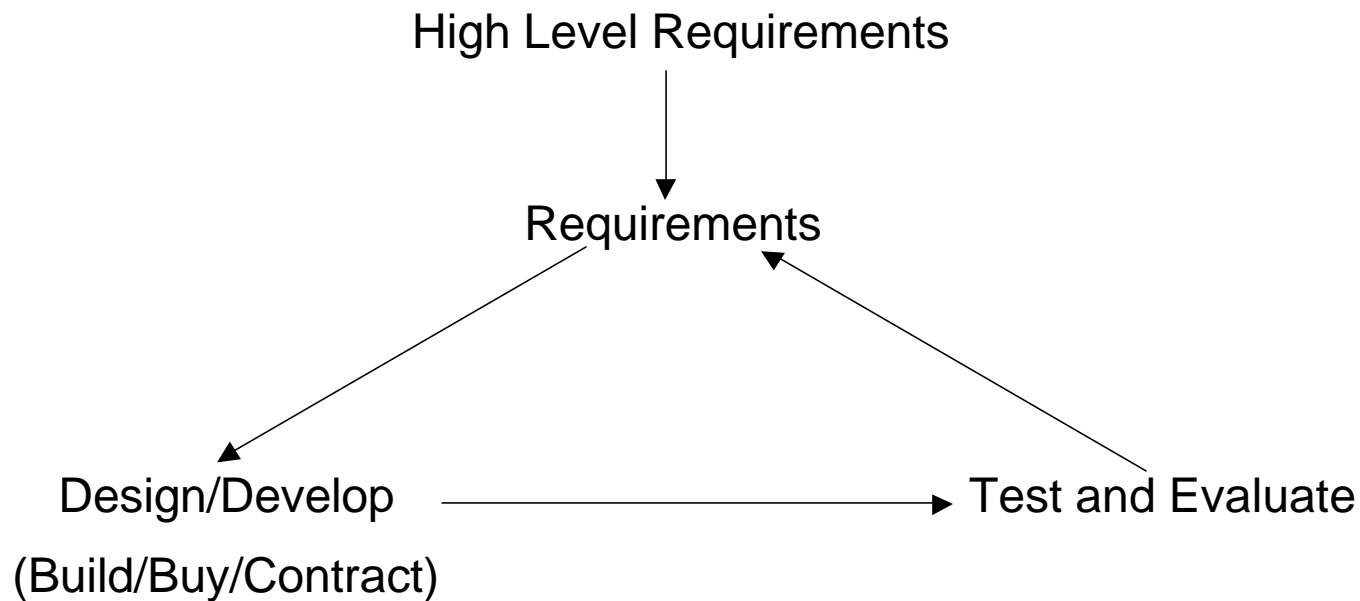
- Decision Domain
 - Problem definition
 - Setting decision context
- System Representation
 - Structuring and mathematically representing the problem
 - Identify data/information flow and analysis strategies
- System Quantification
 - Populating system representation with data/information
 - Estimation/prediction through statistical information integration, including uncertainty quantification
- System Optimization
 - “What if” analyses
 - Uncertainty quantification
 - Sensitivity analysis
- Technology Transfer



Information Integration Framework



Our View of Your Acquisition Process



Model and Simulation Driven

Problem and System Structuring

Deborah Leishman, Ph.D.
leishman@lanl.gov



Overview

- Knowledge Capture and Representation
- A Template for Structured Interviewing
- An Example

Knowledge Capture and Representation

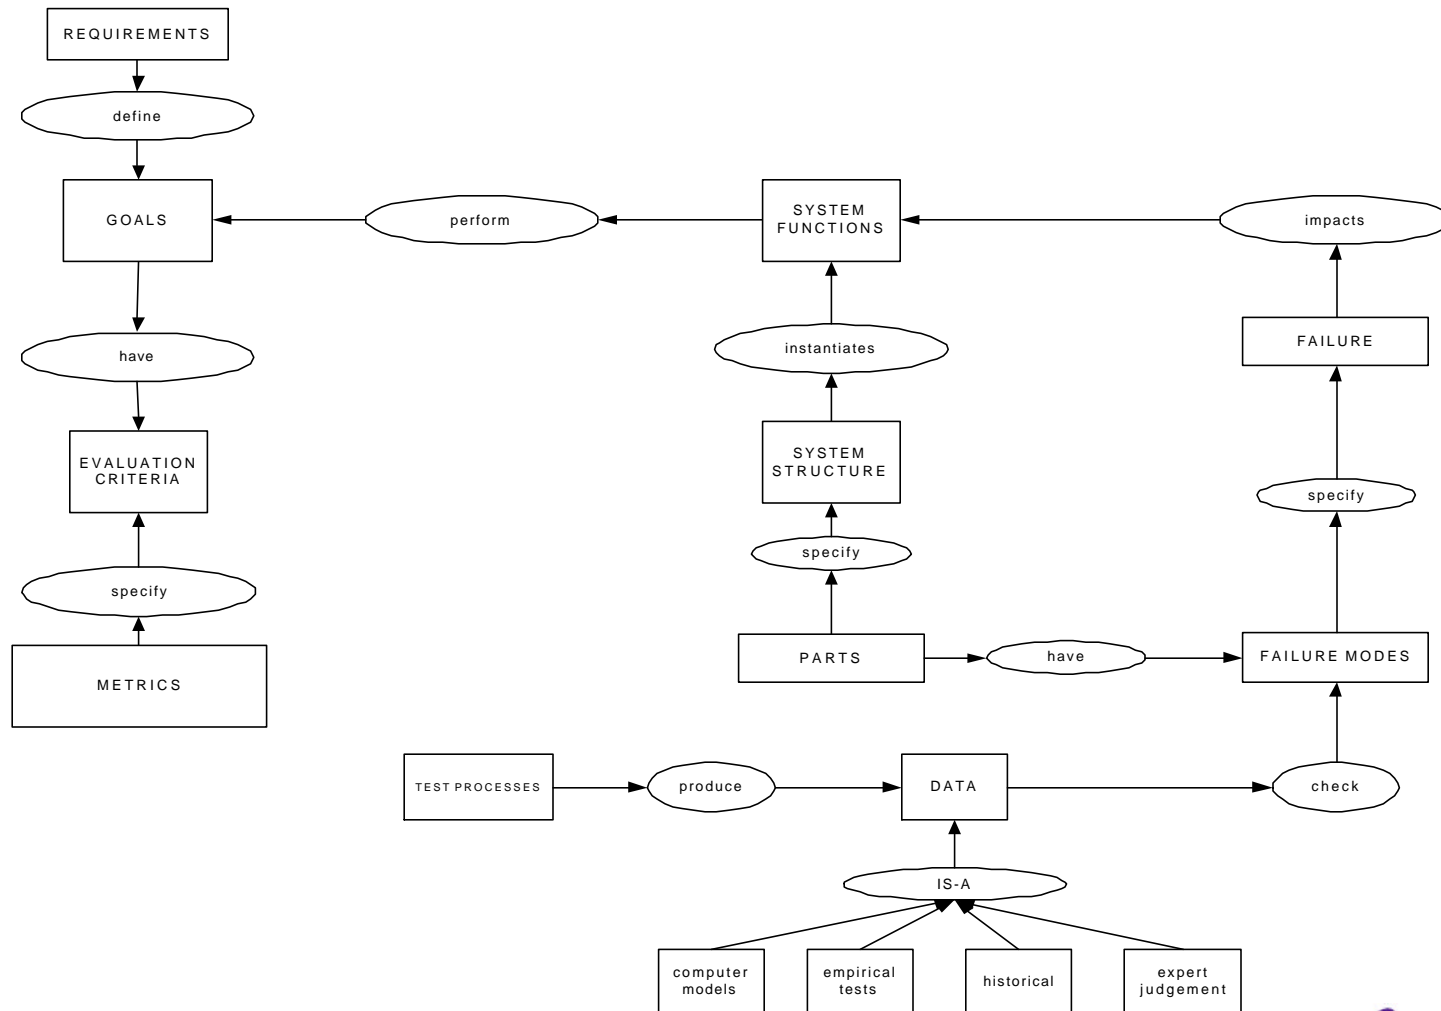
- Knowledge Capture and Representation are done as an iterative refinement process
- Knowledge Capture occurs as a structured interviewing process driven by templates represented as conceptual graphs
- Conceptual Graphs are a Knowledge Representation technique developed by John Sowa (www.uah.edu/~delugach/CG/)
- Notes taken during the structured interviews are refined into a conceptual graph representation in stages (portions of the templates done in different interviews)
- The Knowledge Models are transformed into Proto Influence Diagrams and finally into Statistical Models



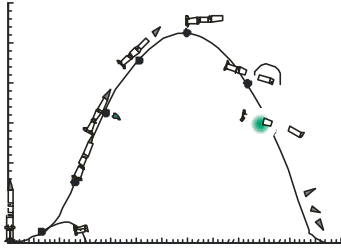
Knowledge Capture and Representation

- First Frame the Problem
 - Problem Definition
 - Decision context
- The CG Knowledge Model template is then used to develop the following descriptions:
 - Decision Goals and Evaluation Criteria
 - System Structure
 - System Functions
 - Test Processes and Failure Modes
 - Data Sources

The Conceptual Graphs KM Template



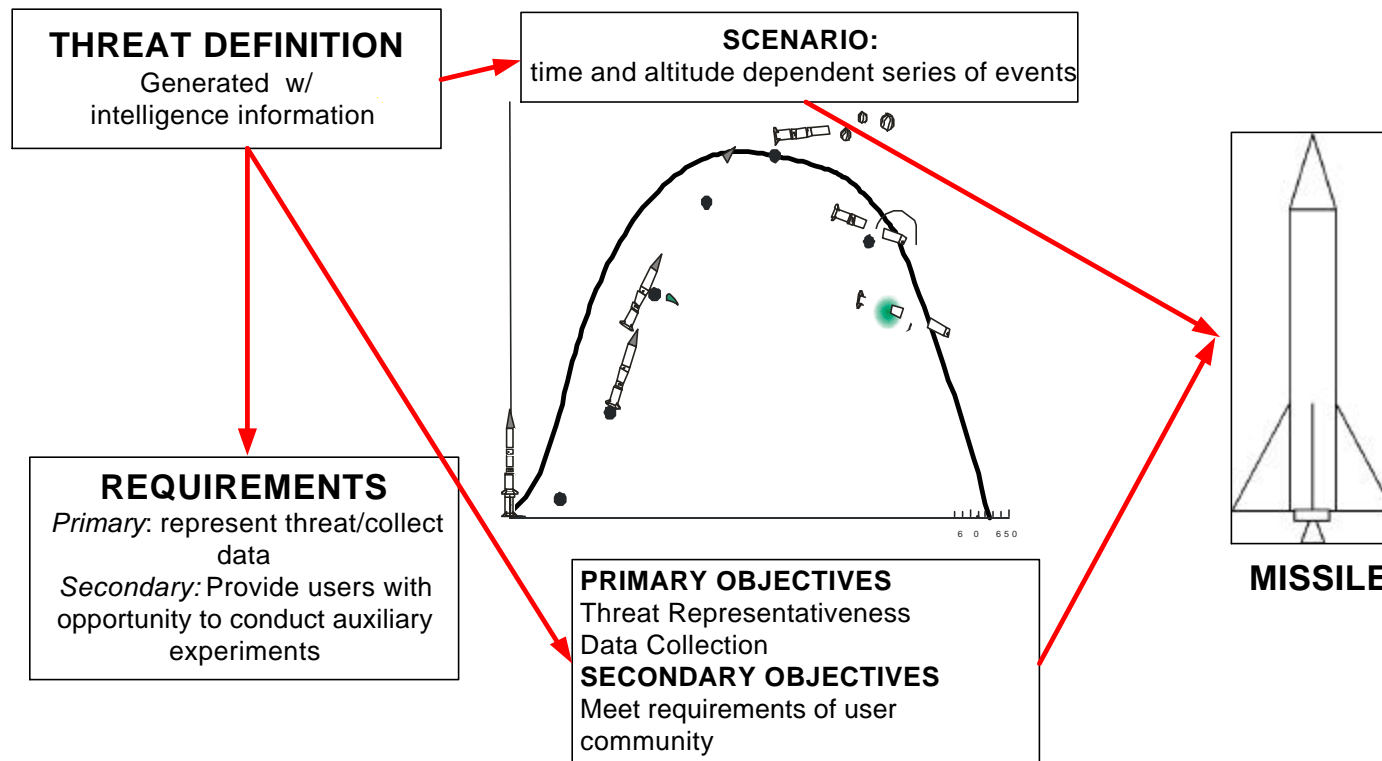
An Example Reliability Model



Missile Defense Agency Critical Measurements Program

- **GOAL:** Fly a high-fidelity, threat-representative missile system for Theater Missile Defense data collection and interoperability exercise
- **ISSUES**
 - Multiple partners and contractors
 - High reliability demanded
 - Full system testing not an option

Problem Definition and Decision Context

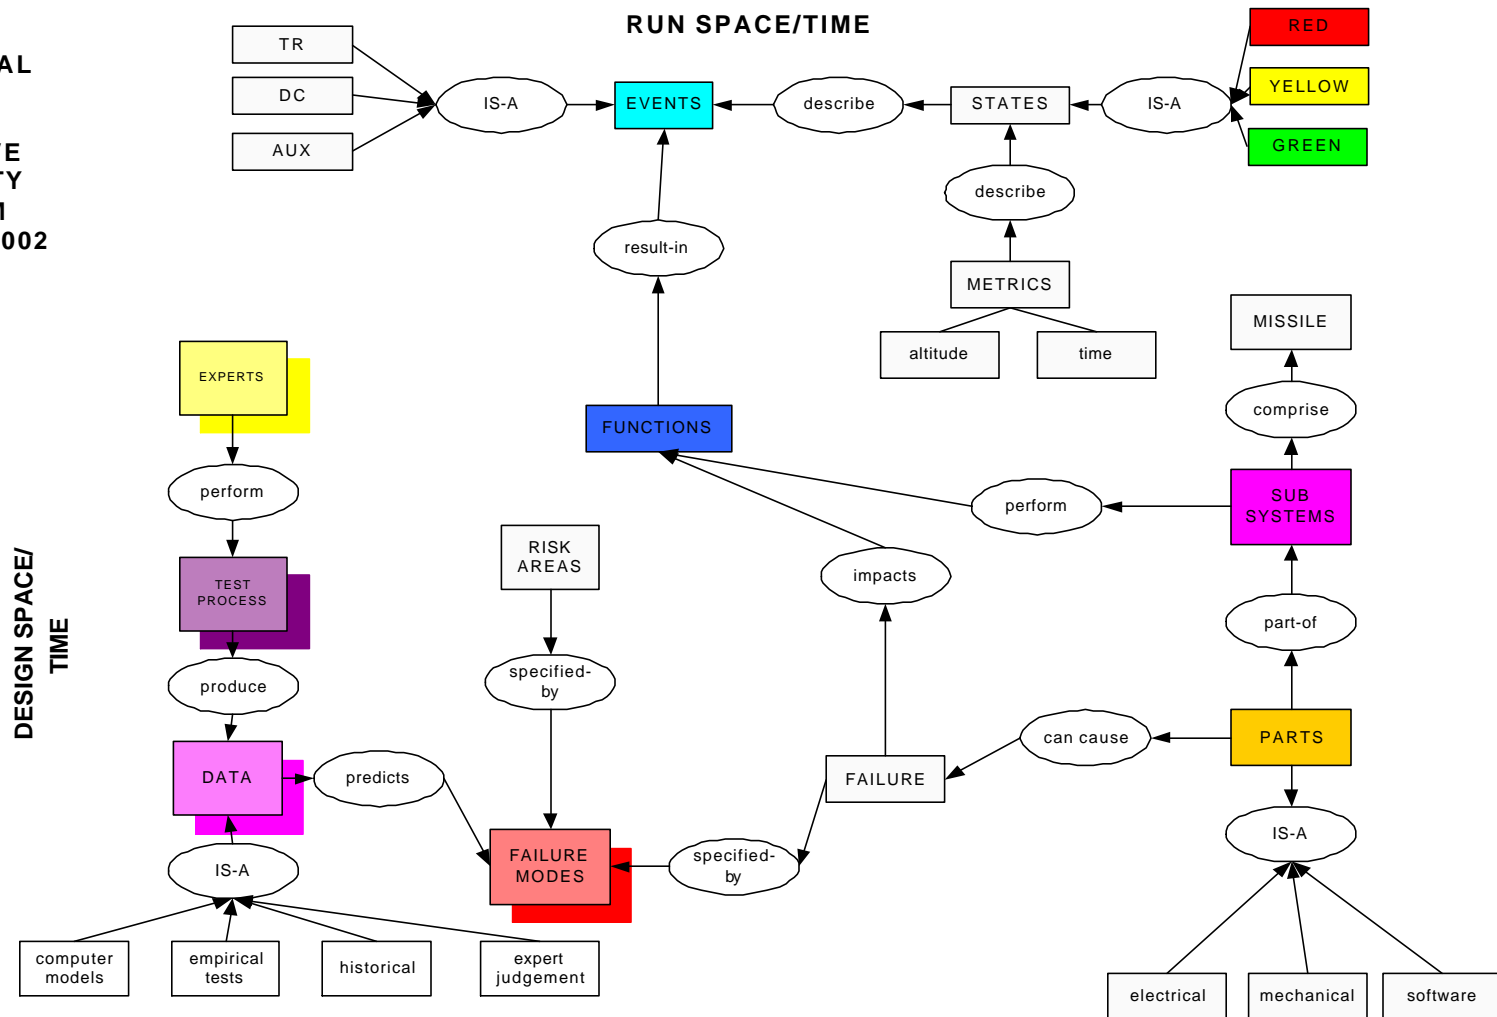


DECISION CONTEXT

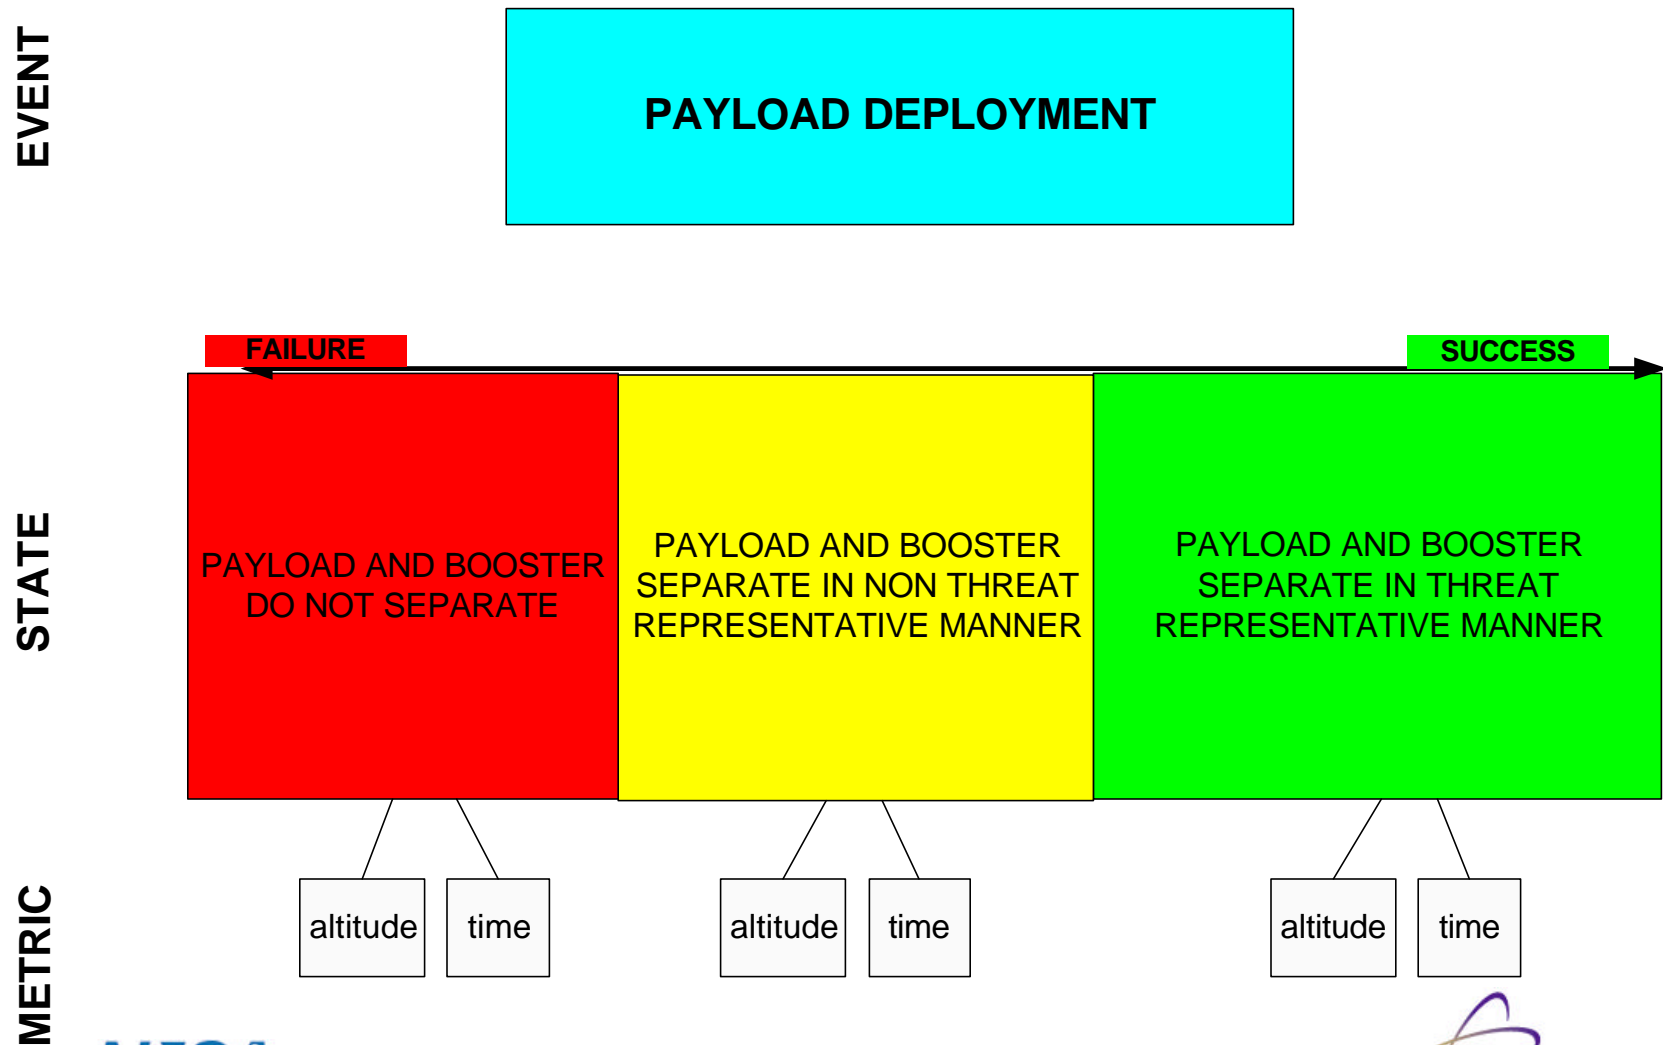
1. Highest risk areas based on reliability estimate
2. Trade-offs between data collection and threat representativeness
3. Prioritization of User Community requirements

First Refinement of the KM Template

CONCEPTUAL
GRAPH
MDA
PREDICTIVE
RELIABILITY
PROBLEM
MARCH 20, 2002

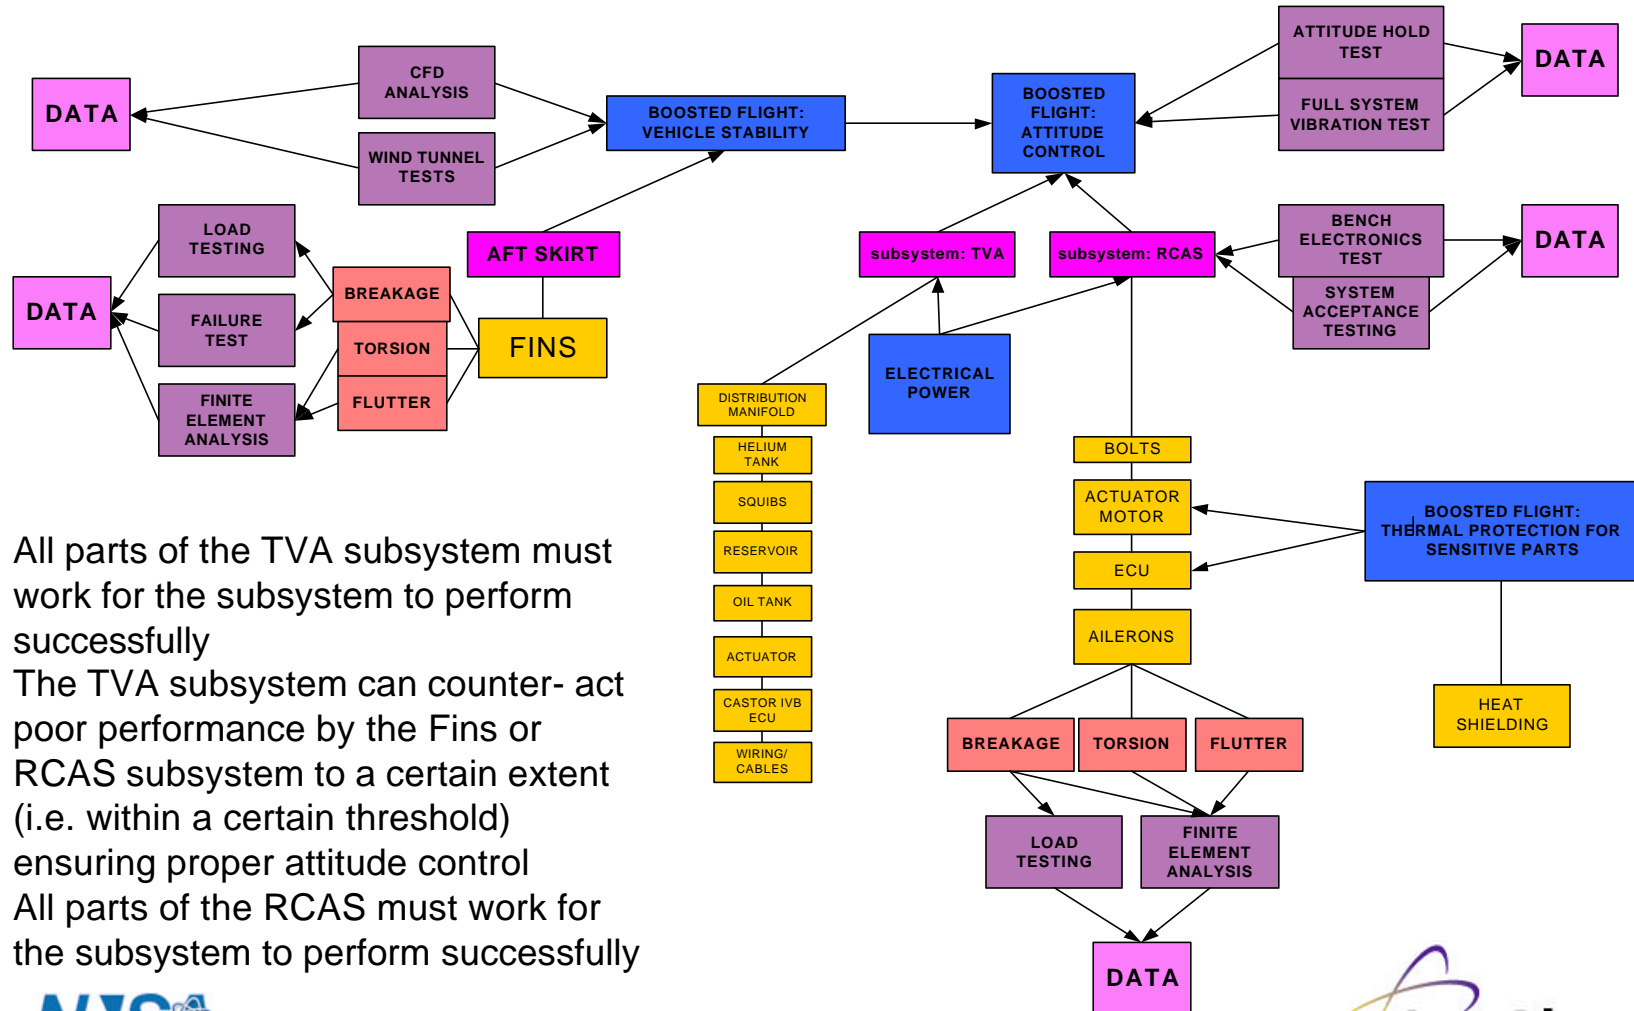


A Further Refinement of the Goals and Evaluation Criteria



The First Proto Influence Diagram

design data for key subsystems: AFT SKIRT AND ATTITUDE CONTROL
payload separation event



- All parts of the TVA subsystem must work for the subsystem to perform successfully
- The TVA subsystem can counter-act poor performance by the Fins or RCAS subsystem to a certain extent (i.e. within a certain threshold) ensuring proper attitude control
- All parts of the RCAS must work for the subsystem to perform successfully



Summary

- Knowledge Capture and Representation happen as an Iterative Refinement Process
- Templates are used to drive a structured Interviewing process
- First Frame the problem
- Refine the Problem by linking Structure, Functions, Test Processes, and Data to Goals

Sources of Uncertainty

Alyson Wilson, Ph.D.
agw@lanl.gov



Qualitative and Quantitative Representations

In the previous section, we discussed two forms of representation: the conceptual graph and the “proto-influence diagram.” The proto-influence diagram actually draws from a collection of statistical representation techniques that include both tree-based and graph-based models.

Trees and Graphs

The basic tree model is the *decision tree*. At each node of a decision tree there is a question or event; arcs coming from each node correspond to the answers to the question or occurrence of an event. A special case of a decision tree is an *event tree*.

Another useful tree model is the *fault tree*, which traces events, using AND and OR gates, that lead to a failure.

Also used are reliability graphs, which capture the physical interconnection of parts, and state-transition graphs, which generalize the decision tree to multiple states.



Statistical Representations

There is important translation that takes place between the “proto-influence diagram” and the actual statistical calculations. The information from the knowledge modeling is transformed into a statistical representation.

“In this way models can be adjusted and elaborated without needing to confront a client with numerical evaluations of uncertainty (e.g., probabilities) early in the analysis—a process about which many clients harbor great suspicion.”

Graphs

When we talk about graphs, we are talking about the formal mathematical kinds of graphs that contain nodes and arcs. The two kinds of graphs that are most commonly used are:

- Reliability block diagrams, where the nodes capture components and functions and their dependencies (series, parallel, k-of-n)
- Graphical models, where the nodes are random variables and the arcs capture conditional dependencies. We are specifically working with *chain graphs*, which are acyclic (no directed cycles) graphs with directed or undirected edges.



Data, Information, and Knowledge

- Expertise
- Expert judgment
- Historical test data
- Data / information on similar or relevant systems
- Design specifications
- Computer simulation model outputs
- Physical model / code outputs
- Test Data

Our goal is to represent the entire state of knowledge about a given problem at a given time.



Expert Judgment as Data

Expert judgment shares traits with data from tests, experiments, or physical observations.

- It is affected by the process of gathering it
- It has uncertainty, which can be characterized and subsequently analyzed.
- It can be conditioned on various factors, such as
 - the phrasing of the question,
 - the information the experts considered,
 - the experts' methods of solving the problem, and
 - the experts' assumptions.
- It can be combined with other information/data.

Quantifying Expert Judgment

Distributions can be formulated by:

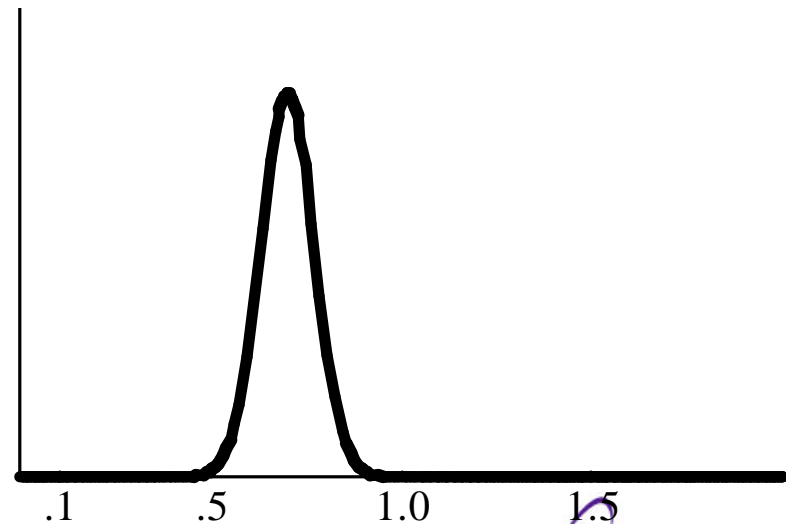
- Having the expert draw a distribution
- Using elicited moments, parameters, or quantiles
- Using elicited membership functions

While we are focusing on using expert judgment to formulate probability distributions, we also have done work using non-probabilistic uncertainty characterizations like Dempster-Shafer theory and fuzzy logic.

Formulating Distributions

Moments—while an expert might be able to estimate a mean, it is extremely rare that he/she would be able to estimate a standard deviation or variance. As such, studies do not recommend this estimation.

Distribution is normal
with a mean of 0.7
and a standard deviation
10% of mean



Formulating Distributions

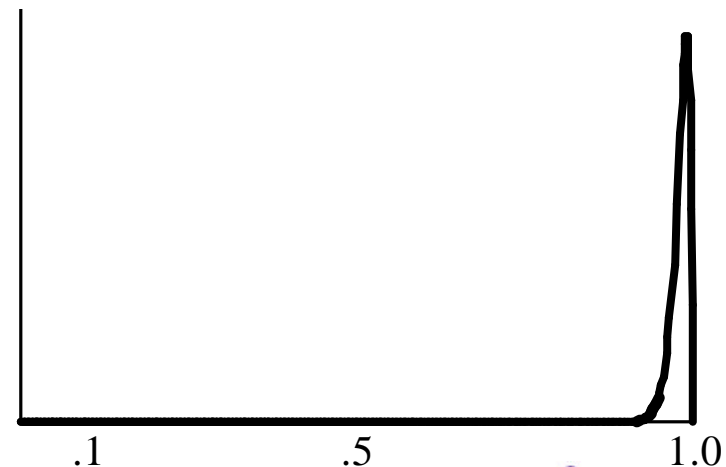
Parameters—rarely can parameters be directly estimated by experts. One such possible case is with distributions whose parameters have interpretations (e.g., 1st beta parameter can be number of successes, and the 2nd parameter can be number of failures).

Beta:

1st = 98 successes

in 100 trials

2nd = 2 failures

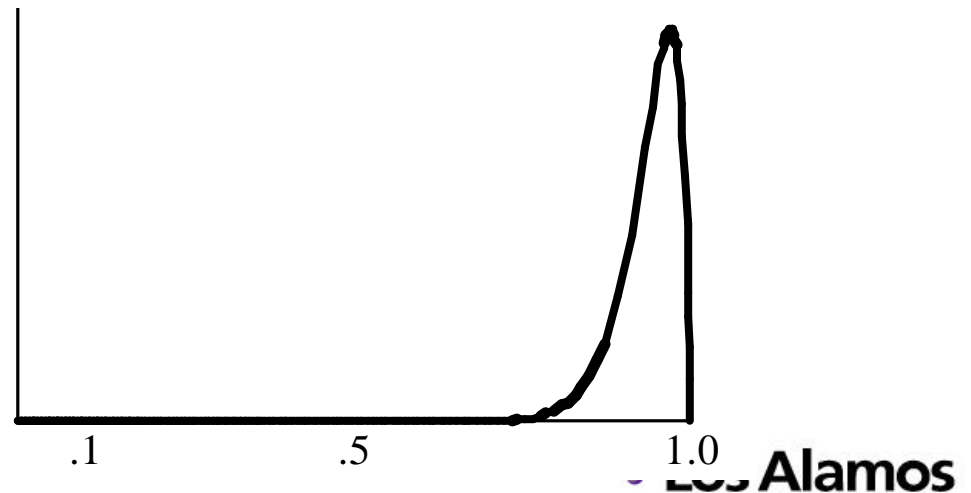


Formulating Distributions

Quantiles—most common. Experts do well in estimating the median as the most likely value or as their best estimate. Studies indicate if an expert provides a mean, it often is a median. Ranges of values (best/worst or max/min) are good for estimating uncertainty; however take into account the experts to underestimation of uncertainty bias.

$$0 \leq p \leq 1$$

$$p_{\max}=0.99, p_{\min}=0.85$$



Which method do you use?

- What kind of expert judgment did you elicit?
- What method is the most tractable?
- What matches the features that your expert considers important?
- Which one can you simulate from?

Computer Model Evaluation

One of the active research areas in our group is the characterization of uncertainties associated with predictions of physically-based computer models.

The statistical research areas include:

- Design of experiments
- Sensitivity/importance analysis
- Feature extraction
- Statistical emulators
- Assessment of model adequacy
- Model calibration
- Extrapolation and prediction



Statistical Framework

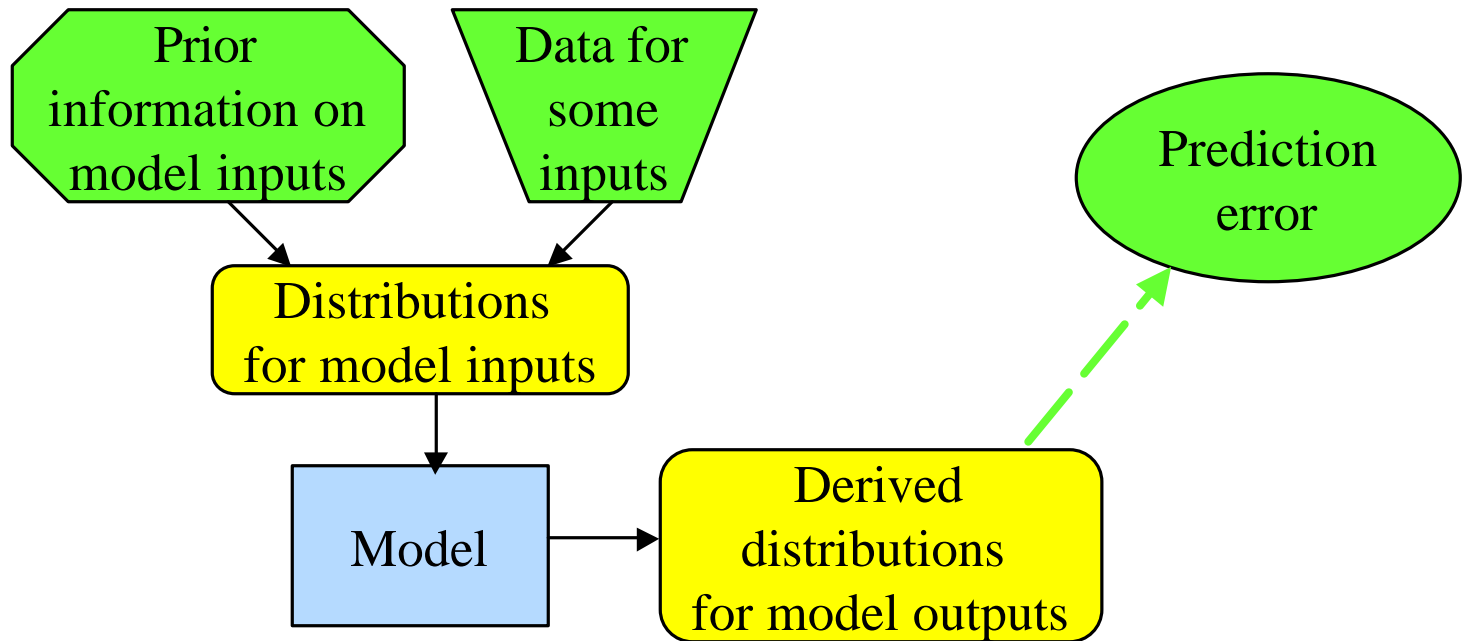
Data y collected under scenario x_{test} are related to model M with parameters θ by

$$y(x_{\text{test}}) = M(x_{\text{test}}; \theta) + b_y(x_{\text{test}}) + e$$

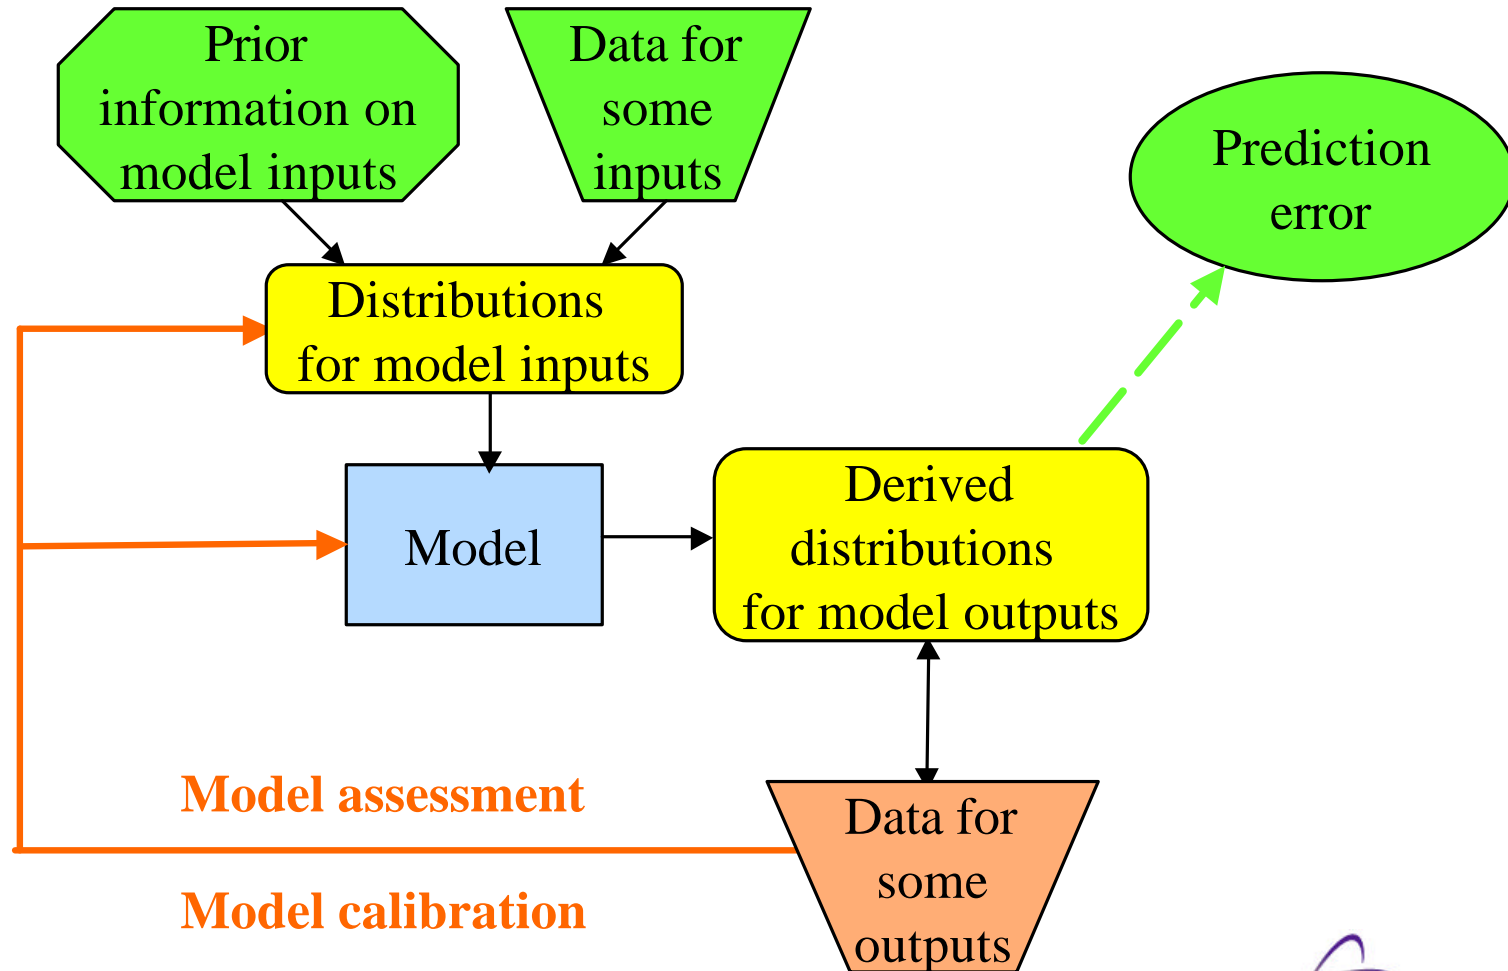
Predictions of z under scenario x_{pred} will be estimated using the model M by

$$z(x_{\text{pred}}) = M(x_{\text{pred}}; \theta) + b_z(x_{\text{pred}})$$

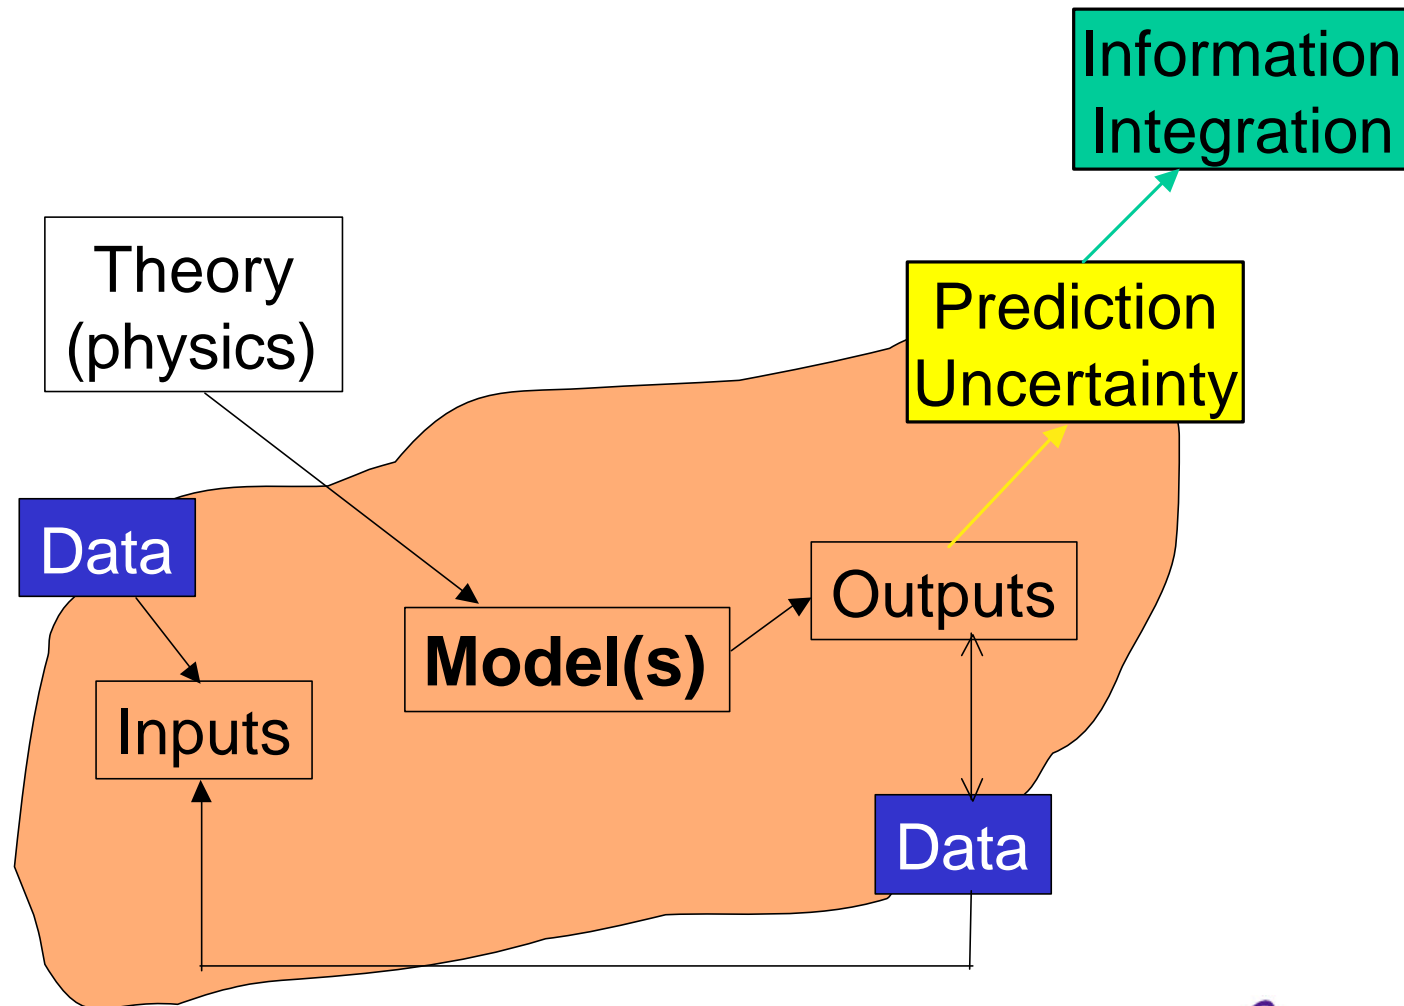
The Forward Problem



Using Data on Output Variables



Statistical Context for CME



Historical and Similar System Data

Historical and similar system data are becoming recognized as important sources of information for system evaluation. In a T&E context, think about the push to use developmental test data to inform the operational evaluation.

The current state of the art is *statistical modeling*, examples of which are given in the next section.

Bayesian Hierarchical Modeling

C. Shane Reese
reese@statmail.byu.edu



Outline

- (brief) Literature review
- What is a hierarchical model
- Conceptual example
- Real data example
- Other applications
- Conclusions

Literature Review

- Draper *et al.* (1992) is the best overview of HM
- Robbins (1955) is the first demonstration of HM, he calls it *empirical Bayes*
- Efron and Morris (1975) also refer to HM as *empirical Bayes*
- Lindley and Smith (1972) and Smith (1973) present a general hierarchical linear model
- (1998 – present) a **flood** of papers on hierarchical models

What is a HM?

- Consists of three parts

1. Observational model: distribution of the data

$$(y_i | \theta_i) \sim f_i(y_i | \theta_i), \quad i = 1, \dots, n$$

(independently)

2. Structural model: distribution of unobservables (parameters)

$$(\theta_i | \alpha) \sim g(\theta_i | \alpha), \quad i = 1, \dots, n$$

(independently)

3. Hyperparameter model: distribution of parameters from (2)

$$\alpha \sim h(\alpha)$$

- How does this relate to regular Bayesian methods?



What is a HM?

- Standard Bayesian Methods

$$p(\theta | y) = \frac{f(y|\theta)g(\theta)}{\int_{\theta} f(y|\theta)g(\theta)d\theta}$$

- Now (under HM):

$$p(\theta | y) = \frac{\prod_{i=1}^n f(y_i|\theta_i)g(\theta_i|\alpha)h(\alpha)}{\int_{\theta_1} \cdots \int_{\theta_n} \prod_{i=1}^n f(y_i|\theta_i)g(\theta_i|\alpha)h(\alpha)d\theta}$$

- Difficulty:

1. Under standard Bayesian methods, calculating

$$m(y) = \int_{\theta} f(y|\theta)g(\theta)d\theta$$

2. Now (under HM), calculating

$$\int_{\theta_1} \cdots \int_{\theta_n} m(y) = \prod_{i=1}^n \int_{\theta_i} f(y_i|\theta_i)g(\theta_i|\alpha)h(\alpha)d\theta$$



Computation (not covered in this course)

- If $m(y)$ is known (that is if the posterior distribution is a known form), then calculation is EASY (calculators).
- If $m(y)$ is not known, then we have to use fancy computational techniques called Markov Chain Monte Carlo (MCMC).

Conceptual Example

Suppose we have data on k components which are believed to have similar (but not the same!!!) reliability.

- Observational model:

$$f(y_i | \theta_i) = \binom{n_i}{y_i} \theta_i^{y_i} (1 - \theta_i)^{n_i - y_i}$$

for $i = 1, \dots, k$.

- Structural model

$$g(\theta_i | \alpha) = \frac{\Gamma(\tau + \zeta)}{\Gamma(\tau)\Gamma(\zeta)} \theta_i^{\tau-1} (1 - \theta_i)^{\zeta-1}$$

that is, $\theta_i \sim \text{Beta}(\tau, \zeta)$ where $\alpha = (\tau, \zeta)$ and

Conceptual Example

- Hyperparameter model

$$h(\tau, \zeta) = \frac{1}{b_{\tau}^{a_{\tau}} \Gamma(a_{\tau})} \tau^{a_{\tau}-1} \exp\{-\tau/b_{\tau}\} \times \\ \frac{1}{b_{\zeta}^{a_{\zeta}} \Gamma(a_{\zeta})} \zeta^{a_{\zeta}-1} \exp\{-\zeta/b_{\zeta}\}$$

that is, τ and ζ have gamma distributions with parameters

(a_{τ}, b_{τ}) , and (a_{ζ}, b_{ζ}) , respectively.

Real Data Example

An anti-aircraft missile has several components (names omitted to protect the innocent).

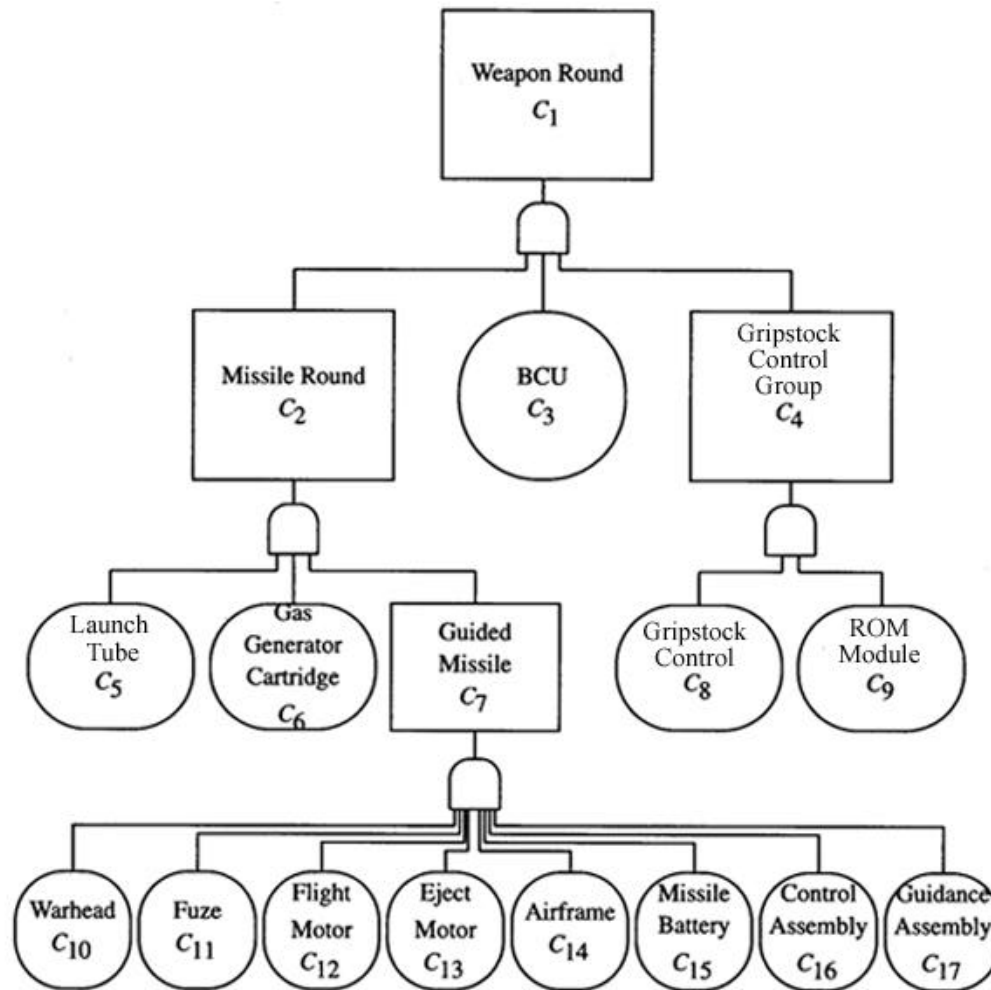
- number of successful tests at a component is binomial
- each component has its own reliability
- Y_i is number of successful tests out of n_i tests
- π_i is the probability of success at each component (reliability)
- Y_i has a Binomial distribution with parameters n_i and π_i

Real Data Example

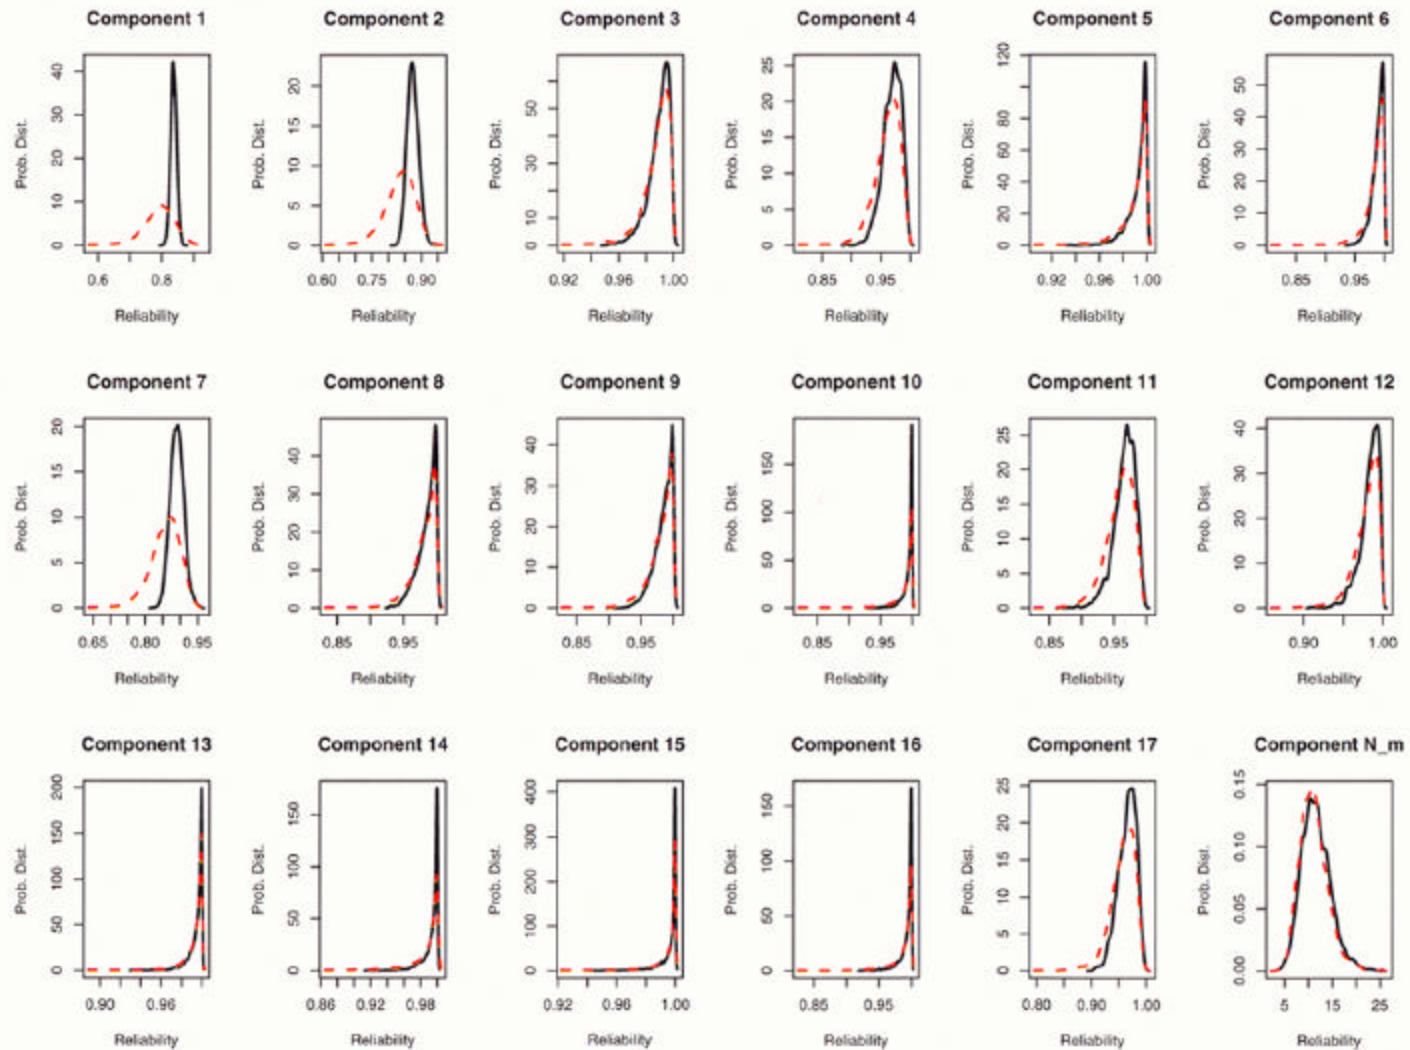
Data

- Over 1400 full system tests (at highest level)
- 45 tests at SOME of the components
- 126 tests at one subsystem
- some components/subsystems have 0 tests (no data)
- Picture

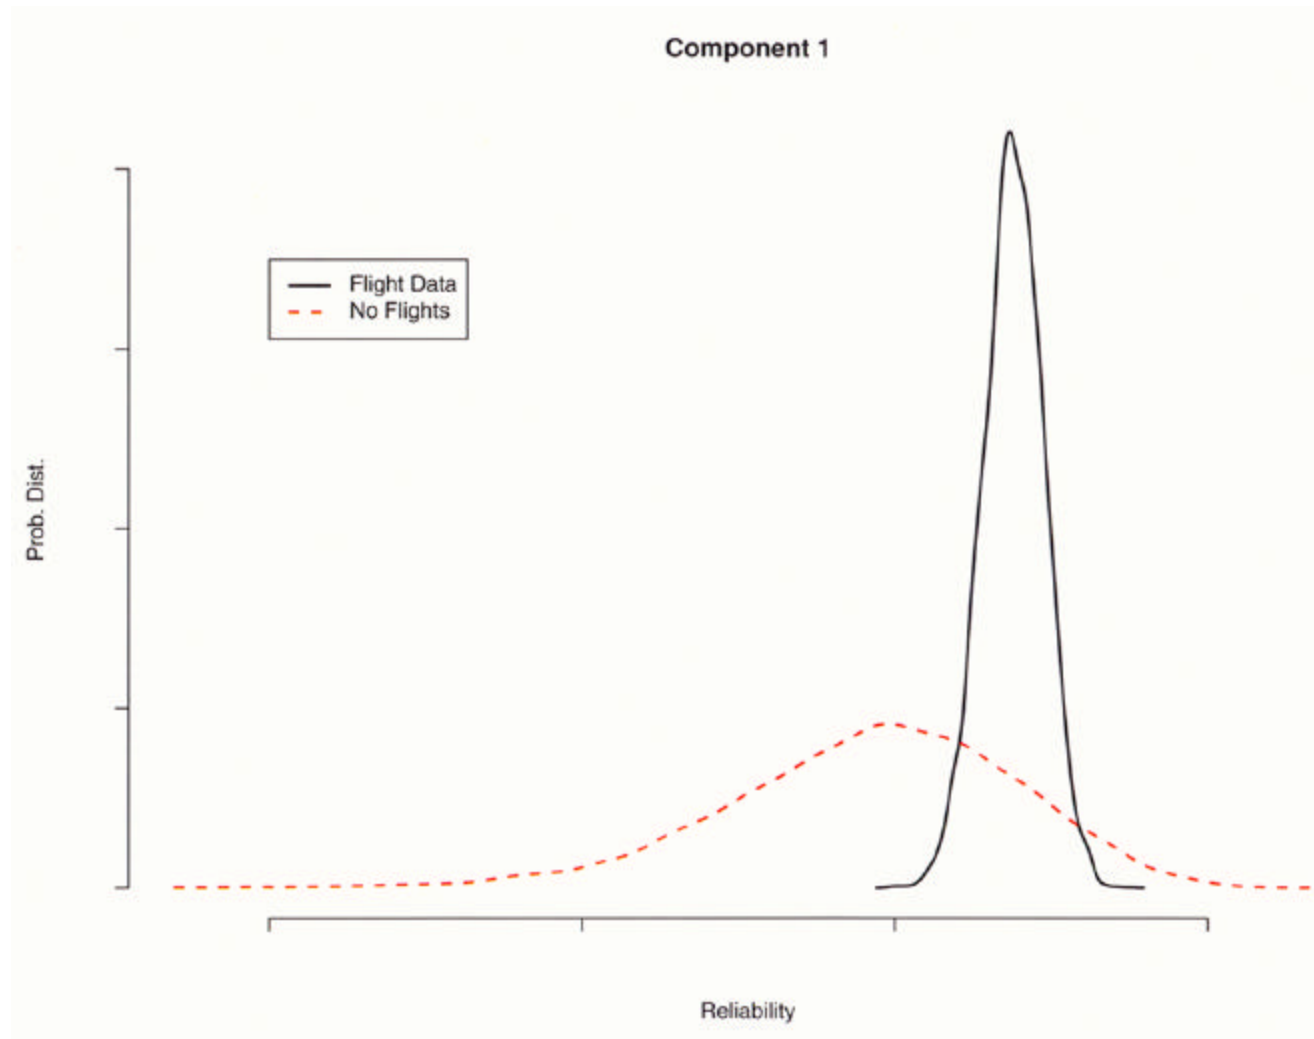
Real Data Example



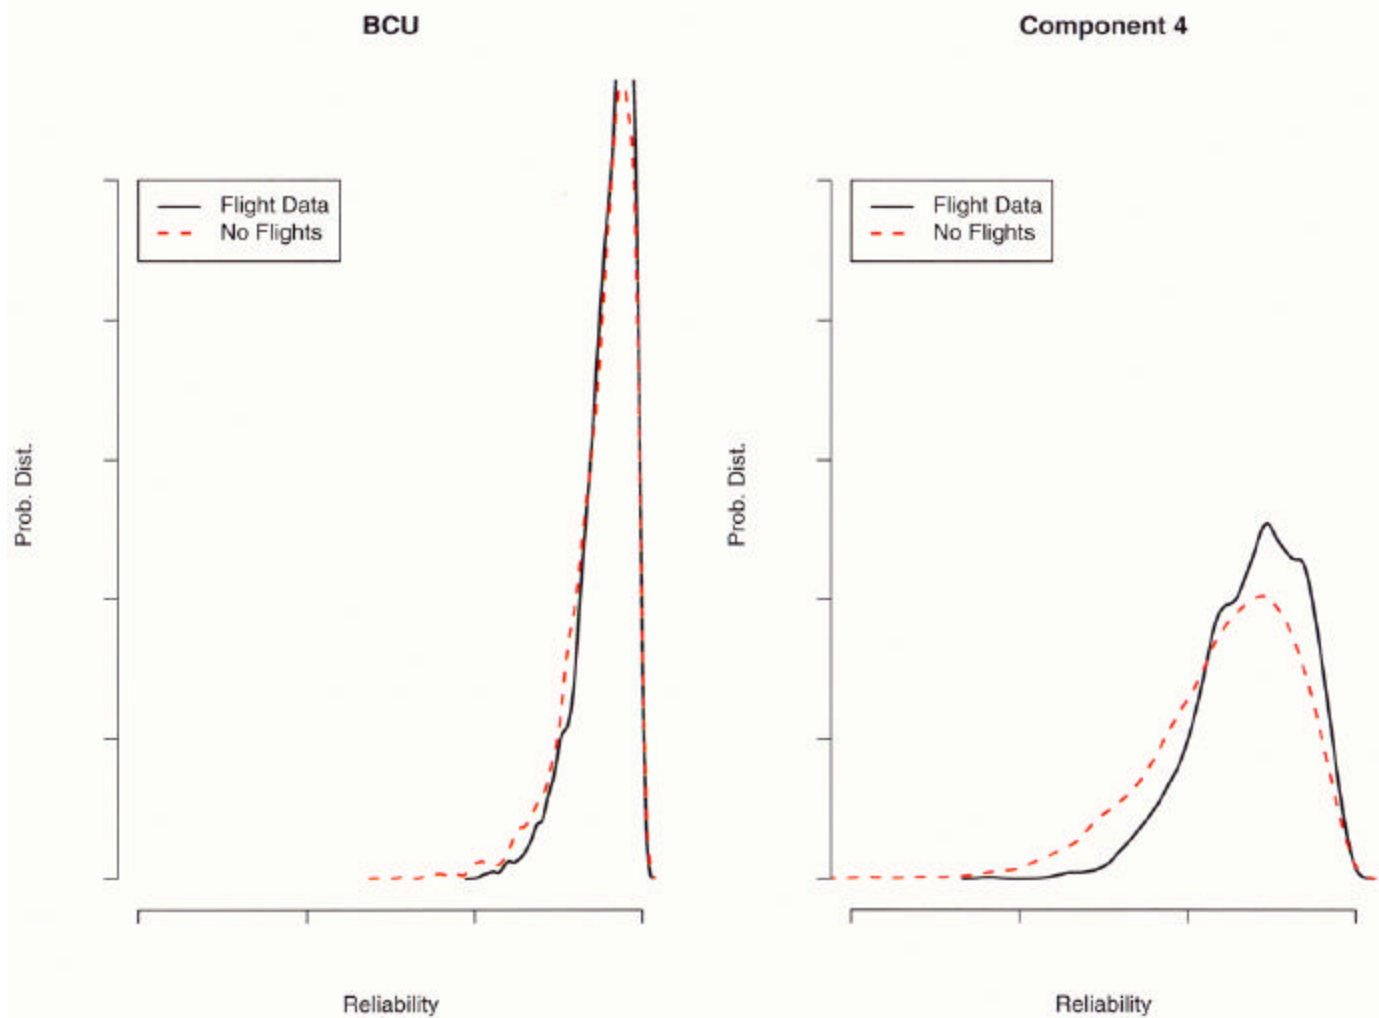
Real Data Example



Real Data Example



Real Data Example



Other Applications

- Bayesian Hierarchical Models are useful for combining diverse information
- Computer Models (biases)
- Physical Experiments (correct, but expensive)
- Historical Data (“prior” information)
- Expert Opinion (“prior” information)
- Reese, Wilson, Martz, Hamada, Ryan (2002) treats case of combining Computer Experiments with Physical Experiments and Expert Opinion using Bayesian Hierarchical Models
- Data sources are different but intended to answer same problem

Background

- Computer/Physical experimental data
- Same (or a subset of the same) factors, but possibly different factor values
- Different responses – transfer function
- Expert opinion
- Simultaneously analyze the combined data using *recursive Bayesian hierarchical model* (RBHM)

Motivation

- Why bother? What do we gain?
 1. More precisely estimated model
 2. Validation of computer experiments
 3. Better predictions
- Cost savings (design?)

Motivation

- The RBHM recognizes important differences between different data sources (expert opinion, computer model, and physical data).
 1. Both location and scale biases in computer models (see Uncertainty and Reliability), allowed to be different for each run of the computer model.
 2. Both location and scale biases in individual experts, allowed to be different for each expert opinion (same or different experts).

Model

- Stage 1
 - Define initial priors on all unknown parameters, including the biases.
 - Update these priors using the expert opinions to form the posterior distributions (using Bayes theorem).

Model

- Stage 2
 - Use the posteriors from Stage 1 as the priors at Stage 2.
 - Update these priors using the computer model output to form new posterior distributions (again by Bayes theorem).

Model

- Stage 3
 - Use the posteriors from Stage 2 as the priors at Stage 3.
 - Update these priors using the physical experimental data to form new posterior distributions (Bayes theorem).
 - This yields the fully updated or final posterior distributions of interest (e.g., regression coefficients, or parameters of a reliability distribution).

Discussion

- We can assess the effect of each data source by comparing the posteriors as they evolve from Stages 1 to 3 (this will be illustrated in the example).
- RBHM can be applied in a linear model framework as well as a reliability context. We will illustrate it in a linear model framework.

Model Details

- Physical experimental data
 - $\underline{Y}_p \sim N(X\underline{\beta}, \sigma^2 I)$, where the physical data \underline{Y}_p are normally distributed with mean $X\underline{\beta}$, X is a model matrix of factors values, and $\underline{\beta}$ is a vector of unknown regression parameters. The notation $\sigma^2 I$ indicates that each physical observation is independent of the others and has variance σ^2 .

Model Details

- Goal
 - The primary goal is to estimate $\underline{\beta}$ and σ^2 and make inferences about them; namely, which components of $\underline{\beta}$ are non-zero or “significant”
 - More appropriately, we want to know which covariates affect the performance metric.

Model Details

- Computer experimental data
 - Comes from complex computer models of physical phenomena, e.g., finite element models.
 - $\underline{Y}_c \sim N(X\underline{\beta} + \underline{\delta}_c, \sigma^2 \Sigma_c)$, where $\underline{\delta}_c$ is a vector of model run specific location biases and Σ_c is a matrix of scale biases (again computer model run specific)
 - Usually

$$\Sigma_c = \begin{pmatrix} 1/k_{c_1} & 0 & \cdots & 0 \\ 0 & 1/k_{c_2} & 0 & \cdots \\ \vdots & 0 & \ddots & \cdots \\ 0 & \cdots & \cdots & 1/k_{c_c} \end{pmatrix}$$

Model Details

- Expert opinion data (expert judgment)
 - $\underline{Y}_o \sim N(X\underline{\beta} + \underline{\delta}_o, \sigma^2 \Sigma_o)$, where $\underline{\delta}_o$ is a vector of possible location biases and Σ_o is a matrix of possible scale biases.
 - Usually

$$\Sigma_o = \begin{pmatrix} 1/k_{o_1} & 0 & \cdots & 0 \\ 0 & 1/k_{o_2} & 0 & \cdots \\ \vdots & 0 & \ddots & \cdots \\ 0 & \cdots & \cdots & 1/k_{o_E} \end{pmatrix}.$$

Biases

- How do these biases arise?
 - Location bias: an expert's average value is often either higher or lower than the true mean.
 - Scale bias: when an expert provides, say, a 0.90 quantile on the true response, this elicited value is often in reality a 0.60 or 0.70 quantile (over-valuation of information)

Elicited Space

- How are these expert opinions elicited?
 - An expected response, y_o .
 - A quantile q_ξ for a prespecified probability ξ (e.g., $\xi = 0.9$, and thus the expert believes that 90% of the responses will be below q_ξ).
 - The “worth” of the expert opinion, m_o .

Worth?

- What is meant by the worth of expert opinion?
 - The corresponding number of physical experimental observations equivalent to the opinion.
 - May be fractional (e.g., may be less than 1)
 - Uncertainty about m_o is expressed through a prior distribution, which is then marginalized (integrated out) when applying the *RBHM*.

Computation

- MCMC methods to simulate observations from the posterior distribution.
- Our method uses Gibbs sampling which involves simulation from complete (or full) conditional distributions.
 - Distribution of each parameter conditional on all other parameters and the data
 - When the complete conditional can't be found in closed form, we simulate from the complete conditional distribution using the Metropolis-Hastings algorithm.

Prior Distributions

- Analogous structure for computer model
- Prior distributions

$$\underline{\beta}|\sigma^2 \sim N(\underline{\mu}_o, \sigma^2 C_o)$$

$$\sigma^2 \sim IG(\alpha_o, \gamma_o),$$

$$m_{o_i} \sim Uniform(0.5_{o_i}^{(e)}, 2.0m_{o_i}^{(e)})$$

$$\delta_{o_i} \stackrel{iid}{\sim} N(\theta_o, \xi_o^2)$$

$$k_{o_i} \stackrel{iid}{\sim} G(\phi_o, \omega_o)$$

Prior Distributions

- Hyperprior

- For $\underline{\delta}_o$:

$$\theta_o \sim N(m_{\theta_o}, s_{\theta_o}^2)$$

$$\xi_o^2 \sim IG(a_{\xi_o^2}, b_{\xi_o^2})$$

- For \underline{k}_o

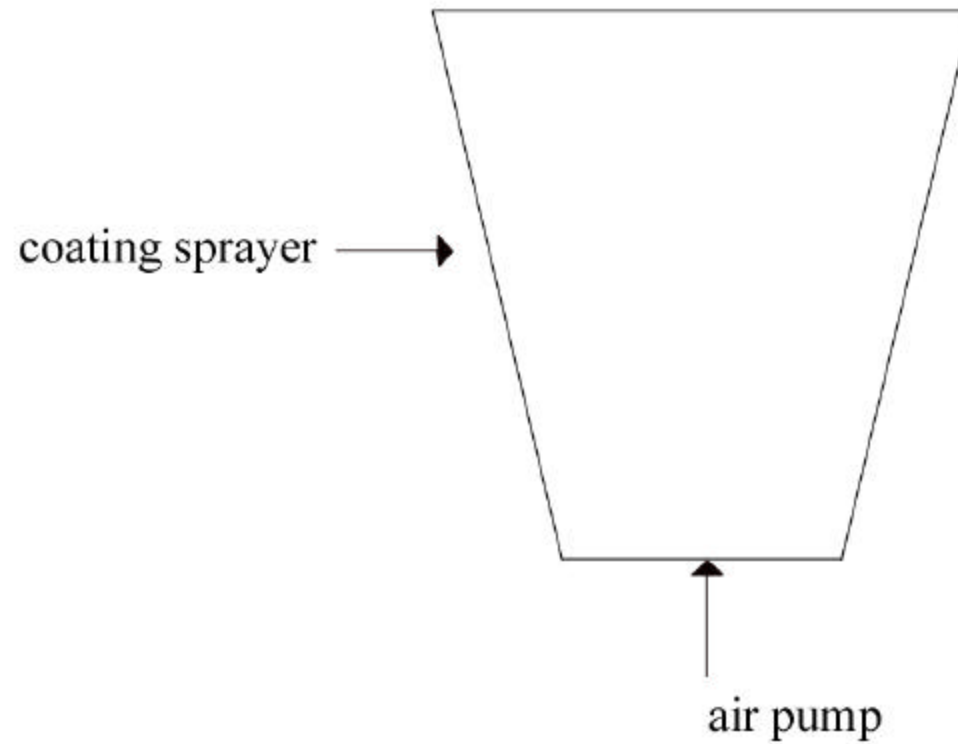
$$\phi_o \sim G(a_{\phi_o}, b_{\phi_o})$$

$$\omega_o \sim G(a_{\omega_o}, b_{\omega_o}),$$

Example

- Fluidized Beds used to coat food products
- Air is used to “float” the product through for even coating

Example



Example

- Three thermodynamic computer models (with increasing fidelity) were developed.
- Response: Steady-state thermodynamic operating point (Y)
- Input variables:
 - Pump air temperature (A)
 - Fluid velocity (V)
 - Coating solution flow rate (R)
 - Atomization air pressure (P)
 - Room Humidity (H)
 - Room temperature (T)

Example

- 28 runs of each computer model (at different combinations of input variables) for a total of 28 x 3 computer model runs.
- 28 runs of the physical machine at each of the combinations of input variables.
- There are differences between “data” sources

Example

- Model

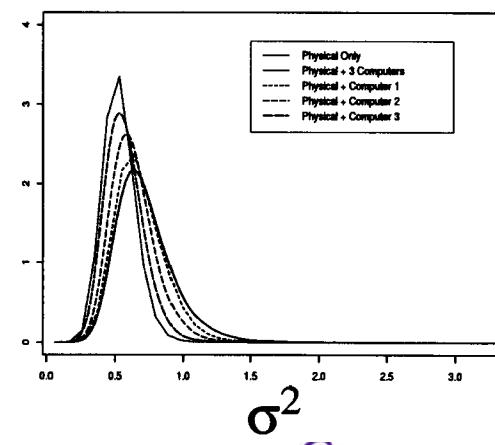
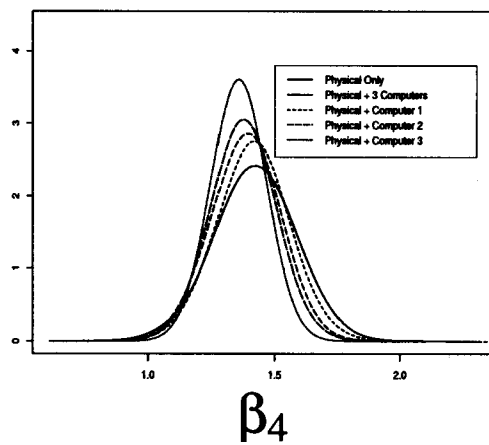
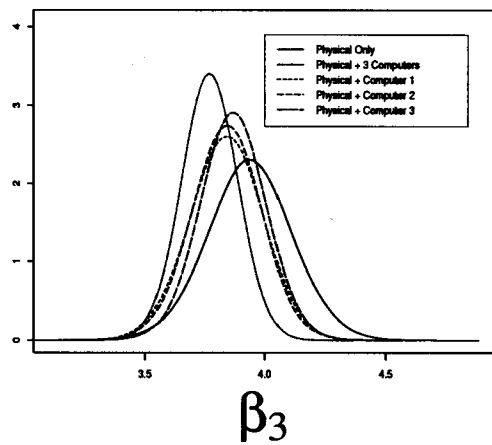
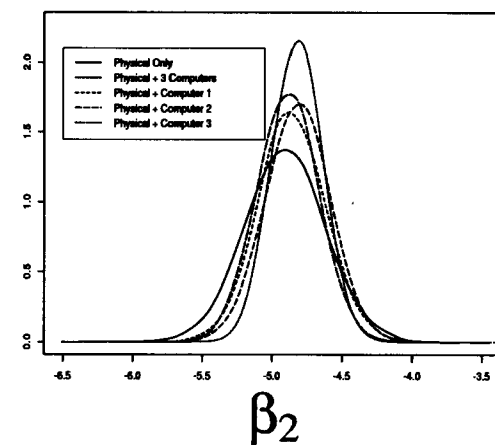
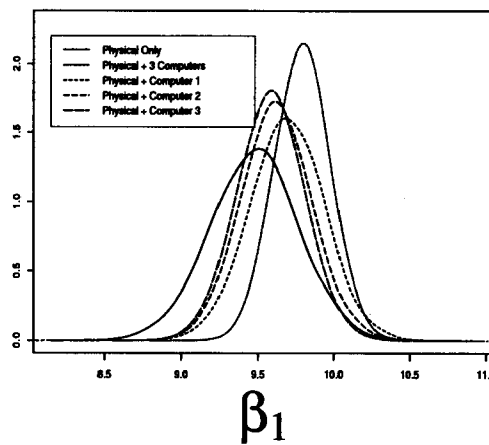
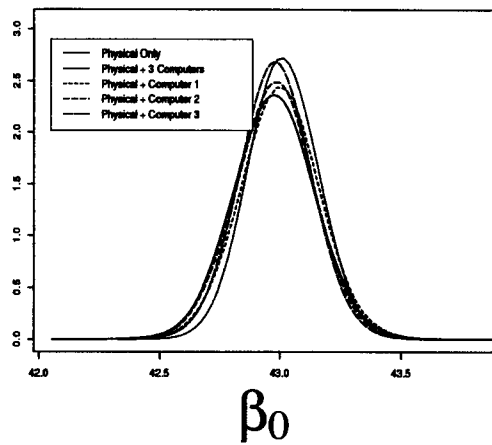
$$\begin{aligned} E(Y_p) &= \underline{X} \underline{\beta} \\ &= \beta_0 + \beta_1 A + \beta_2 R + \beta_3 V + \beta_4 (R \times V) \end{aligned}$$

and

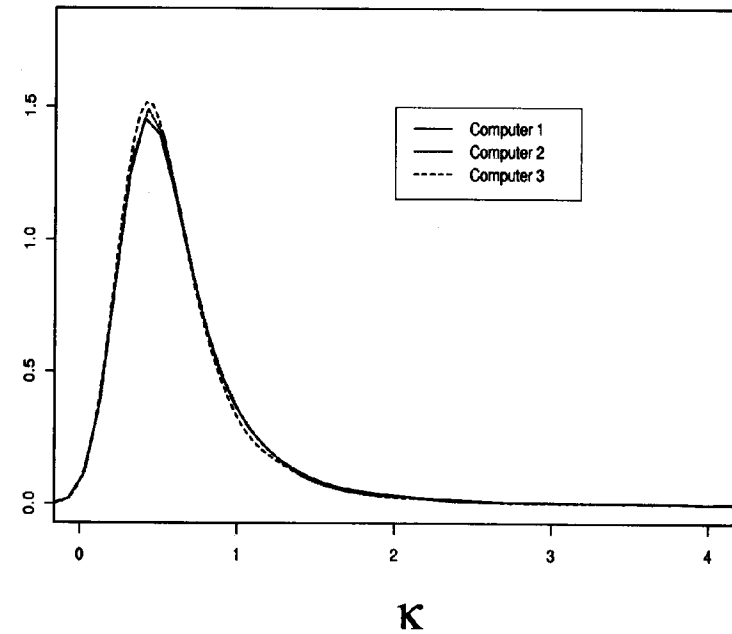
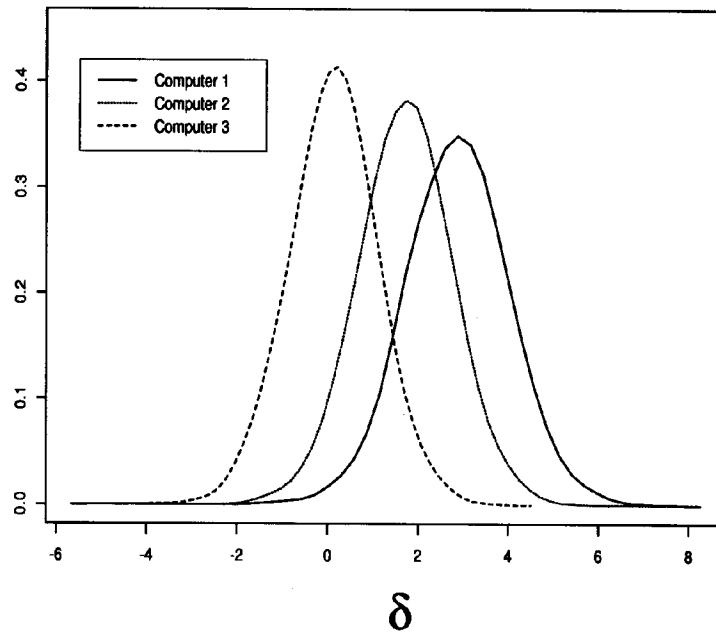
$$V(Y_p) \equiv \sigma^2$$

- Goal: Estimate $\underline{\beta} = (\beta_0, \beta_1, \beta_2, \beta_3, \beta_4)$ and σ^2 .

Example



Example



Discussion

- More precise estimation of parameters
- Predictive distribution of biases provides validation of computer models
- Wide applicability
 - Example is for performance metrics in linear models framework
 - Reliability distributions are minor modification
- Complicated models can be handled

Optimal Allocation of Resources

C. Shane Reese
reese@statmail.byu.edu



Design

- Want the most bang for your buck (literally)
- Total cost (TC) of computer and physical experiments:
 - $TC = FC_c \times I\{n_c = 1\} + n_c C_c + FC_p \times I\{n_p = 1\} + n_p C_p$, where the indicator function $I\{\bullet\} = 1$ if its argument is true and 0 otherwise

where FC_c is the “start-up” cost for a computer experiment, $I\{n_c = 1\}$ is the indicator that *some* computer experiments will be run, n_c is the number of computer experiment runs, C_c is the cost of each computer run, where FC_p is the “start-up” cost for a physical experiment, $I\{n_p = 1\}$ is the indicator that *some* physical experiments will be run, n_p is the number of physical experiment runs, C_p is the cost of each physical experiment run.

Optimal Design

- Let $U(D_c, D_p)$ be a measure of the amount of information in a combined design with computational experiment D_c and physical experiment D_p .
- The experimental design problem then becomes:
 1. Find the combined design (D_c, D_p) that maximizes $U(D_c, D_p)$
 2. subject to the constraint $TC \leq B$ (where B is the budget).
- Could be time as well as cost!



Optimal Design

- Choices that must be made:
 1. Must specify (possibly with uncertainty)
 FC_c, C_c, FC_p, C_p
 2. Must choose a form for $U(D_c, D_p)$
 3. Bayesian: maximize the expected utility

$$E[U(X)] = \int \int U(x) d\mathbf{q} dy$$

where θ is the unknown parameters, y is the data, and X is the design matrix.

4. Popular choices:
 - Shannon information
 - Determinant of $(X'X)^{-1}$ (D-optimality)



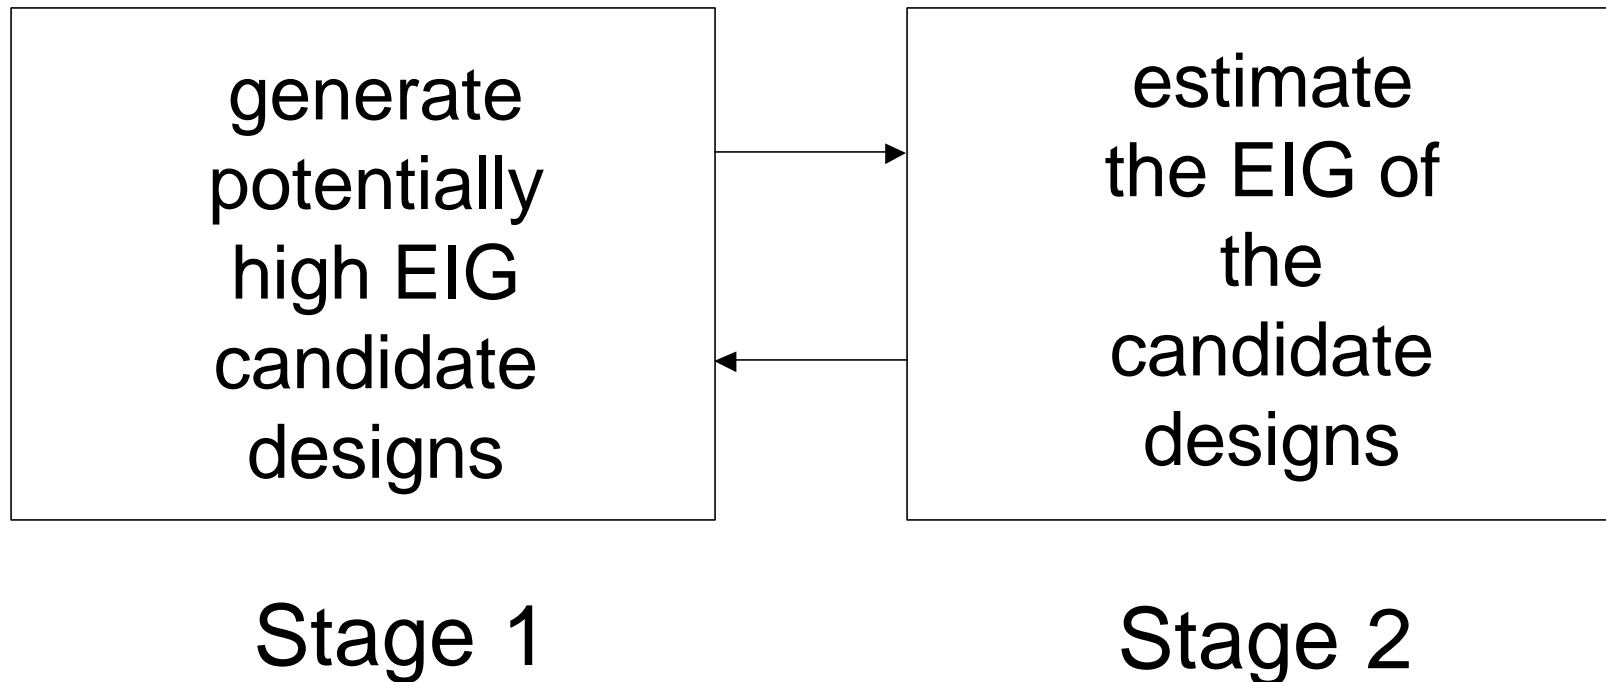
Choices of Utility

- We choose: gain in Shannon information
- Mathematically:

$$U(x) = \int \int \log[\mathbf{p}(\mathbf{q} \mid y, X)] f(y \mid \mathbf{q}, X) \mathbf{p}(\mathbf{q}) d\mathbf{q} dy$$

- Hard to calculate
- Even harder to maximize

Algorithm



Two-stage Iterative Bayesian Experimental Design Solver



How are we going to maximize?

- Genetic algorithm
 1. Choose an initial set of designs
 2. Allow crossovers
 3. Allow mutations
- Multiple generations
- We never get “total” optimality, but we get VERY close
- Stochastic optimization

Conclusions

- Uncertainty is attached to every problem
- Inclusion of expert judgment
- Good allocation of costs
- Considers the importance/value of each information source

Decision Analysis

Alyson Wilson, Ph.D.
agw@lanl.gov



What are the parts of a decision analysis?

1. A set of available actions or decisions; one action must be selected
2. Uncertain states of nature that impact what the consequences are for each decision
3. For each action/state pair, a utility

Example

You get to choose which weapon you want to take into a room full of bad guys. If you have more bullets than bad guys, and the weapon works, you win. Otherwise, you lose.

Two choices (actions):

a_1 = 95% reliable weapon with 8 bullets

a_2 = 60% reliable weapon with 13 bullets

There are somewhere between 0 and 19 bad guys (states of nature). Each count has a 5% chance.

Example

Decision analysis says to choose the action that has the highest expected utility. Let W denote the utility of winning and L denote the utility of losing.

$$\begin{aligned} E[U(a_1)] &= 0.05W + 0.4(0.95W + 0.05L) + 0.55L \\ &= 0.43W + 0.57L \end{aligned}$$

$$\begin{aligned} E[U(a_2)] &= 0.05W + 0.65(0.6W + 0.4L) + 0.3L \\ &= 0.44W + 0.56L \end{aligned}$$

As long as $W > L$ (you get more for winning than losing), the expected utility of a_2 is larger, so a_2 is the correct decision.

Value of Information

Suppose that you could pay to find out for sure whether there are ten or fewer bad guys in the room or more than ten bad guys in the room. How much is that worth to you?

Calculate the expected utility for each case:

$$\begin{array}{ll} \leq 10 & \begin{array}{l} E[U(a_1)] = 0.78W + 0.22L \\ E[U(a_2)] = 0.64W + 0.36L \end{array} \end{array} \quad \text{Choose } a_1$$

$$\begin{array}{ll} > 10 & \begin{array}{l} E[U(a_1)] = L \\ E[U(a_2)] = 0.2W + 0.8L \end{array} \end{array} \quad \text{Choose } a_2$$

$$0.55(0.78W + 0.22L) + 0.45(0.2W + 0.8L) - 0.44W + 0.56L = 0.079(W - L)$$



Notes

- Very nice introduction to decision making can be found in Dennis Lindley's *Making Decisions*, 2nd Edition. ("The book is addressed to business executives, soldiers, politicians, as well as scientists; to anyone who is interested in decision-making and is prepared to take the trouble to follow a reasoned argument.")
- Cost-free information is always expected to be of value.
- Notice that some of our calculations depend on the values of the utilities we specify. Actually defining and quantifying utility requires elicitation and hard work.
- The statistical portion of IIT quantifies the probability distributions associated with the states of nature; the knowledge modeling portion of IIT defines the actions, states of nature, and utilities.

Reuse

Deborah Leishman, Ph.D.
leishman@lanl.gov



Reuse

- Knowledge Reuse
- Software Reuse

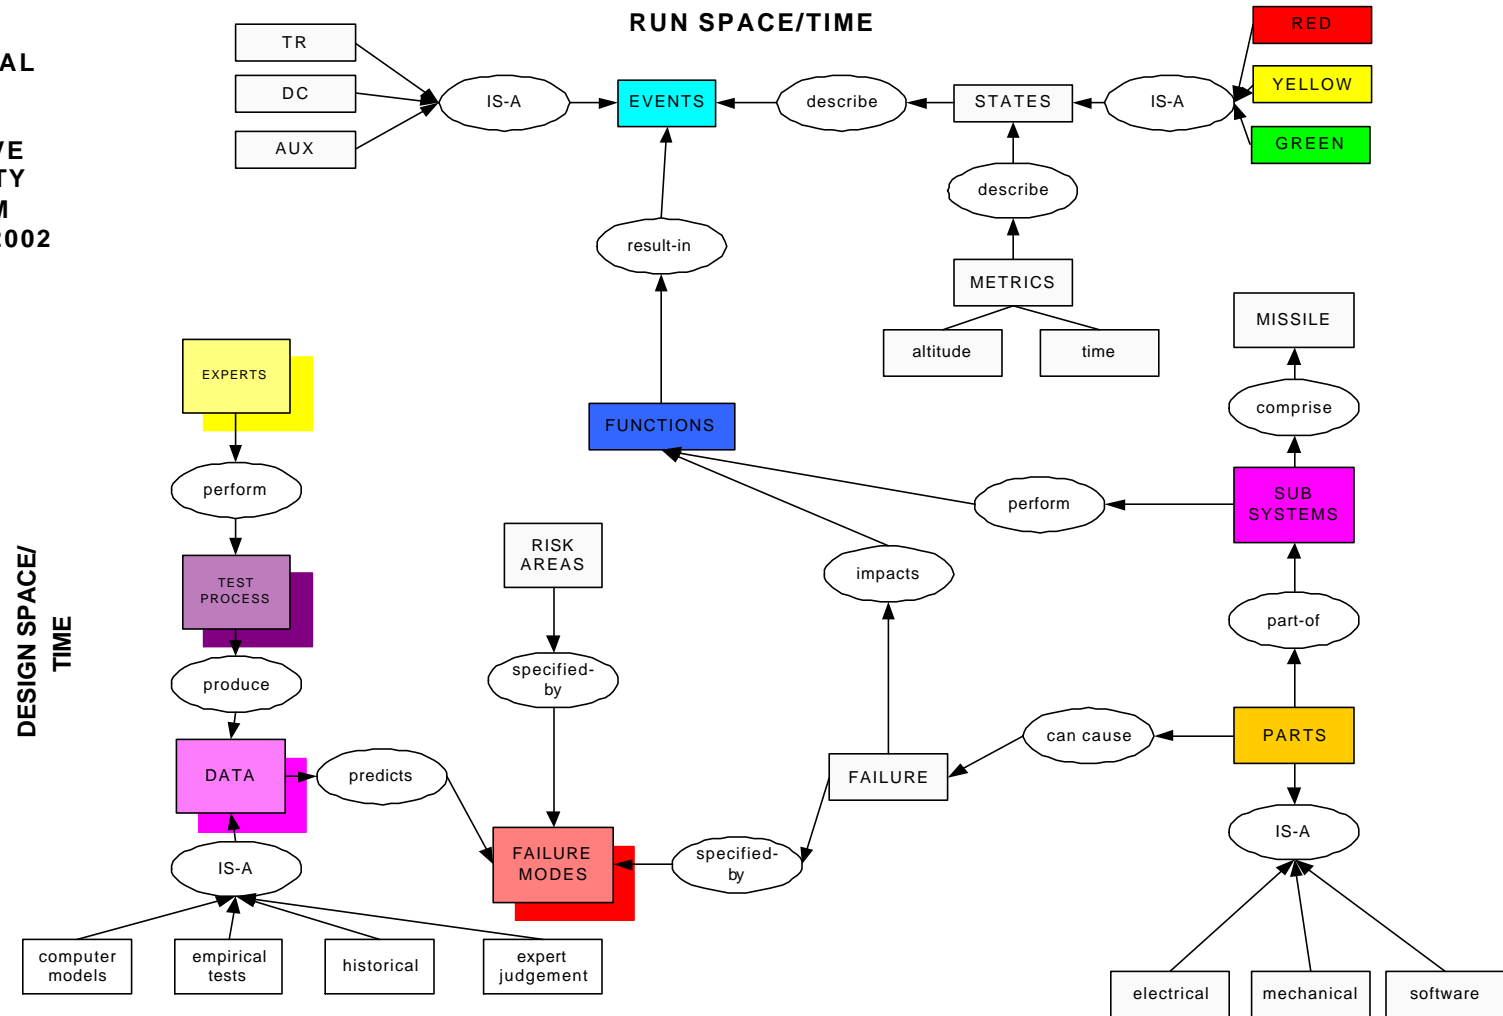
Knowledge Reuse

- Utilize a Knowledge Management tool such as a Lotus Notes Teamroom
- Structure the Teamroom using the CG template's Key Concepts
 - Events, Functions, Subsystems, Parts, Test Processes, Data Sources
- Provides access to more detailed information about the key concepts
- Provides distributed access and security
- Provides a common language for distributed groups of developers
- Provides a common virtual place for teams to meet



First Refinement of the KM Template

CONCEPTUAL
GRAPH
MDA
PREDICTIVE
RELIABILITY
PROBLEM
MARCH 20, 2002



Software Reuse

- Why it hasn't worked for 30 years:
 - Culture – pretty important but we know what to do
 - Tools/Repositories – not really the problem either
 - Technical – Major problem
 - C libraries and such work a bit
 - OO Wrong granularity - its too small
 - Frameworks are too complex
 - Medium sized functional/services based components are just right – we now have what we need
 - The really hard problem will be doing Domain Analysis to define the right sets of cohesive functions/services as components for an industry to evolve
 - c/v (commonality/variability) is where its at



Reuse

For further information: leishman@lanl.gov

